420-TP-010-002

Transition To Release B

Technical paper - Not intended for formal review or Government approval.

Technical Paper

Preliminary

March 1996

Prepared Under Contract NAS5-60000

RESPONSIBLE ENGINEER

Howard Ausden/s/ 3/27/96

Howard Ausden, Senior Scientist/Engineer Date
EOSDIS Core System Project

SUBMITTED BY

Paul Palmer, Development Engineering Manager Date
EOSDIS Core System Project

Hughes Information Technology Systems
Upper Marlboro, Maryland

This page intentionally left blank.

Abstract

This document presents the current state of the analysis of the ECS system Release A-to-Release B transition, discusses the issues that have arisen planning the transition, makes trade-off analyses, and ultimately will lead to the documentation of detailed plans for transition at each site. The installation of Release B at the non-Release A sites is also considered.

The objectives, approach, and timescale are presented. The activities at the ECS sites during the transition timeframe are presented, and contentions are examined. The hardware-software mappings for Release A and B are given, and technical issues arising from changes between the releases are discussed. Trade-offs are made, leading to a statement of some candidate solutions. These include the definition of a sustaining engineering release to install part of the Release B software infrastructure into the deployed Release A configuration. Finally, the transition process is outlined for a generic Release A site.

Keywords: transition, Release B, plan, version, topology, migration.

This page intentionally left blank.

Change Information Page-

List of Effective Pages	
Page Number	**Issue**
Title	Preliminary
iii through x	Preliminary
1-1 through 1-4	Preliminary
2-1 through 2-4	Preliminary
3-1 through 3-12	Preliminary
4-1 through 4-10	Preliminary
5-1 through 5-8	Preliminary
6-1 through 6-2	Preliminary
7-1 through 7-34	Preliminary
8-1 through 8-2	Preliminary
9-1 through 9-6	Preliminary
AB-1 through AB-4	Preliminary

Document History			
Document Number	**Status/Issue**	**Publication Date**	**CCR Number**
420-TP-010-001	Draft	February 1996	
420-TP-010-002	Preliminary	March 1996	96-0244

This page intentionally left blank.

Contents-

1. Introduction

2. Transition Timeframe, Objectives, And Approach

3. Activity Schedule And Requirements

4. Release A Operational Configurations

5. Release B Operational Configurations

6. Program and Project Constraints and Assumptions

7. Identification of Transition Issues

8. Candidate Solutions

9. Site Plans

Figures-

Tables

Abbreviations and Acronyms

1. Introduction

1.1 Purpose

This paper describes the analysis undertaken to date for transitioning ECS to Release B, as referenced at the Release B Incremental Design Review and the subsequent Operations Workshop. The analysis is still at an early stage and will continue through the Release B Consent To Ship Review (CSR-B).

Transition is the process of installing, testing, and making operational the Release B custom software, COTS, and hardware.

The scope of the plan includes:

1. Transition to Release B at sites that have Release A.
2. Transition aspects of the installation of Release B at sites that do not have Release A, such as when and how these B-only sites interoperate with sites that are transitioning from A.
3. FOS is not in the scope of this study, but CSMS upgrades for FOS and integration with the EOC are in scope.
4. Support for Release B, plus ongoing support for Release A, at the ECS Development Facility (EDF) at Hughes in Upper Marlboro.

1.2 Organization

This paper is organized as follows:

- Section 1 Introduction - Purpose, organization, and review process for this document. Definition and scope of transition.
- Section 2 Timescale For Transition, and objectives and approach for transition planning
- Section 3 Activity Schedule And Requirements, presents a schedule for the currently planned transition-phase relevant activities of the many organizations affected by this plan and discusses salient characteristics of each activity shown in the schedule.
- Section 4 Release A Operational Configurations, showing the mapping of custom code and COTS software to processors for release A.
- Section 5 Release B Operational Configurations, showing the mapping of custom code and COTS software to processors for Release B.
- Section 6 Programmatic And Project Constraints And Assumptions - statement of assumptions and conditions for the transition process.

- •≠ Section 7 Identification Of Transition Issues - for each subsystem, a discussion of transition issues, plus a discussion of overall system issues.
- •≠ Section 8 Candidate Solutions - Possible solutions with their merits and demerits, leading to chosen approaches.
- •≠ Section 9 Logistics - a statement of the transition plan. In this first version, one specific possible transition strategy using postulated hardware configurations for one Release B DAAC and one Release A DAAC.

Questions regarding technical information contained within this Paper should be addressed to the following ECS contact:

- •≠ Howard Ausden, Senior Scientist/Engineer, telephone 301-925-1132, e-mail hausden@eos.hitc.com.

Questions concerning distribution or control of this document should be addressed to:

Data Management Office
The ECS Project Office
Hughes Information Technology Systems
1616 McCormick Drive
Upper Marlboro, MD 20774-5372

1.3 Transition Planning Timescale

An update of this paper will be issued in June1996 to incorporate feedback from the April 1996 submittal of this paper and knowledge gained during CDR-B. Transition planning will continue through CSR-B until September 1997, with further submittals of this paper issued as shown in the following table, to document evolving plans for each site.

Table 1-1 Schedule For Transition Plan

420-TP-010-002	April 96	Delivered for Release B CDR. Contains plans for transition at a generic Release A DAAC and at a generic Release B-only DAAC. Emphasis is on technical issues rather than definition of activities at each site, contention for system resources, or site logistics.
420-TP-010-003	June 96	Updated to incorporate feedback received from CDR and comments on 420-TP-010-002. Includes final hardware-software mapping information for Release A and more detailed information for B. Also includes initial assessment of M&O impacts regarding staffing and procedures.
420-TP-010-004	Early 97	Includes final hardware-software mapping information for Release B, detailed discussion of activities at each site, proposed resolutions for system resource contentions, and detailed transition plans for each site.
420-TP-010-005	420-TP-010-004+ 2 months	Incorporates feedback from DAACs et al on individual site transition plans.

1.4 Acknowledgement

This initial submittal of the document is the product of an internal Transition Working Group which includes representatives from each subsystem, M&O, IATO, the Science Office, and the Multi-Release Support group.

Subsequent submittals of the document will incorporate inputs from the external community (ESDIS, DAACs, DAAC liaisons, et al).

This page intentionally left blank.

2. Transition Timeframe, Objectives, And Approach

2.1 Transition Timeframe

An overview of the transition timeframe follows:

1.≠ Conduct Release B walkthroughs (early install and checkout of Release B software at DAACs) May-June 97. The walkthroughs exercise Release B functionality prior to the completion of I&T for pathfinding purposes, to provide early indications of any problems. This is the first time that Release B software is taken to any site.

2. Release B Consent To Ship Review (CSR) June 97

3. Acceptance testing (IATO) for Release B at sites June - August 97.

4. Independent validation and verification (IV&V) for Release B at sites September 97.

5.≠ Transitional activities at Release A sites October - December 97. These include, for example, migration of data from the release A DBMS (Sybase) to the Release B DBMS (Illustra) in the Science Data Server. Also included here are verifications that the new release produces the same results as the old release, when executed on the same data.

6. Release B operational January 98.

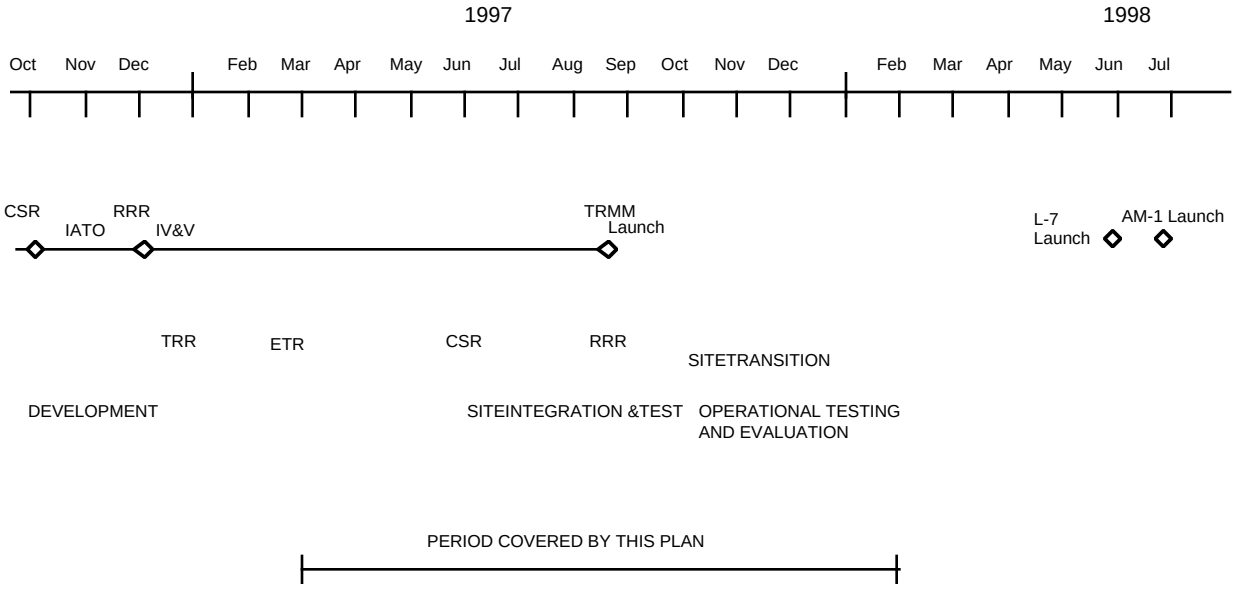

Figure 2-1. The Transition Timeframe

2.2 Objectives

Objectives for the transition of Release A to Release B are:

For Release A, to minimize the perturbation to ongoing operational activities.

For Release B, to ensure the availability of Release A equipment needed to install, configure, test, train, accept, IV&V, and activate Release B.

2.3 Approach

Our analysis takes a three-pronged approach.

The first part of the analysis consists of identifying and resolving technical issues in the transition, as follows:

1.≠ Detailed software to hardware mappings for each site for both Release A and Release B are updated or created, specifying custom executables and COTS versions mapped to each processor.

2.≠ Release A and Release B mappings are compared to identify all COTS, custom code, and hardware upgrades, and identify issues.

3. Candidate issue resolutions are proposed.

4.≠ Trade-off studies are conducted to generate final blueprints for transitional activities at each site.

The second part of our analysis repeats the process used for the IR1 To A Transition Plan. All Release A and early Release B usage requirements are gathered via coordination with DAACs and ESDIS, the equipment and support needs associated with them are assessed, and contention issues are identified. The approach is as follows:

1.≠ An activity schedule for relevant Release A and B activities at each site is developed, including:

- TRMM Operations
- SSI&T
- V0 Migration at sites
- Data Assimilation Office (DAO)
- Mission Readiness Exercises
- ≠ Release A sustaining engineering support in the EDF and at the sites, including COTS installation and test
- Release B development (including EPs) in the EDF
- Release B I&T in the EDF and at sites
- Release B IATO and IV&V in EDF and at sites

- Transitional activities at sites, including folding of Release A science data into Release B
- Activities associated with DAAC-unique interfaces and extensions.

2. The service, hardware, and support requirements for the above are determined, including complexities involved in switching between Release A and B configurations.
3. Issues resulting from contention are derived and evaluated.
4. Candidate solutions (such as early installation of mode management), work-arounds, prioritizations/compromises, and recommendations, and the proposed transition schedule, are developed and analyzed.
5. Cognizant parties are coordinated with, by means of the technical paper process, telecons, and ad hoc meetings.

The third part of the analysis consists of an assessment of the impacts of transition on ECS M&O with regard to staffing, training, ops procedures, ops databases, baseline control, and implementation and management of transition.

This page intentionally left blank.

3. Activity Schedule And Requirements

Release A activities during the transition timeframe are shown in Figure 3-1 and Release B activities are shown in Figure 3-2. The activities are described in the following sections.

Note that this section is preliminary and will be expanded in later submittals of this document (see table 1-1).

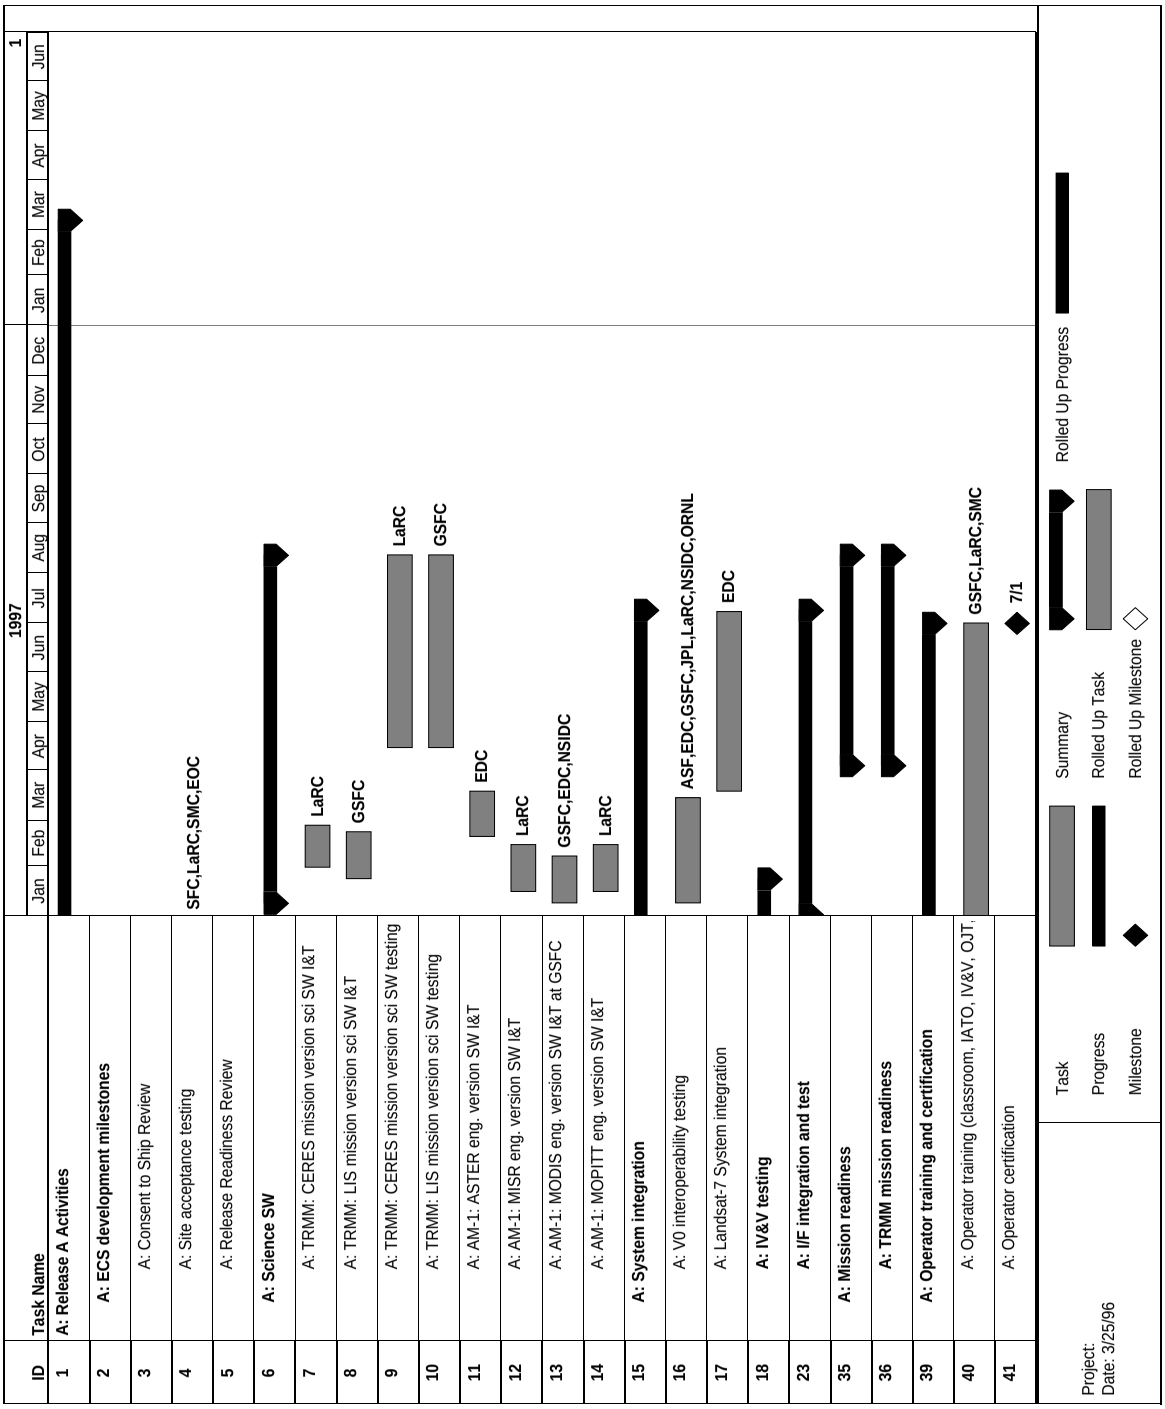

Figure 3-1. Release A Activities During Transition to B (1 of 2)

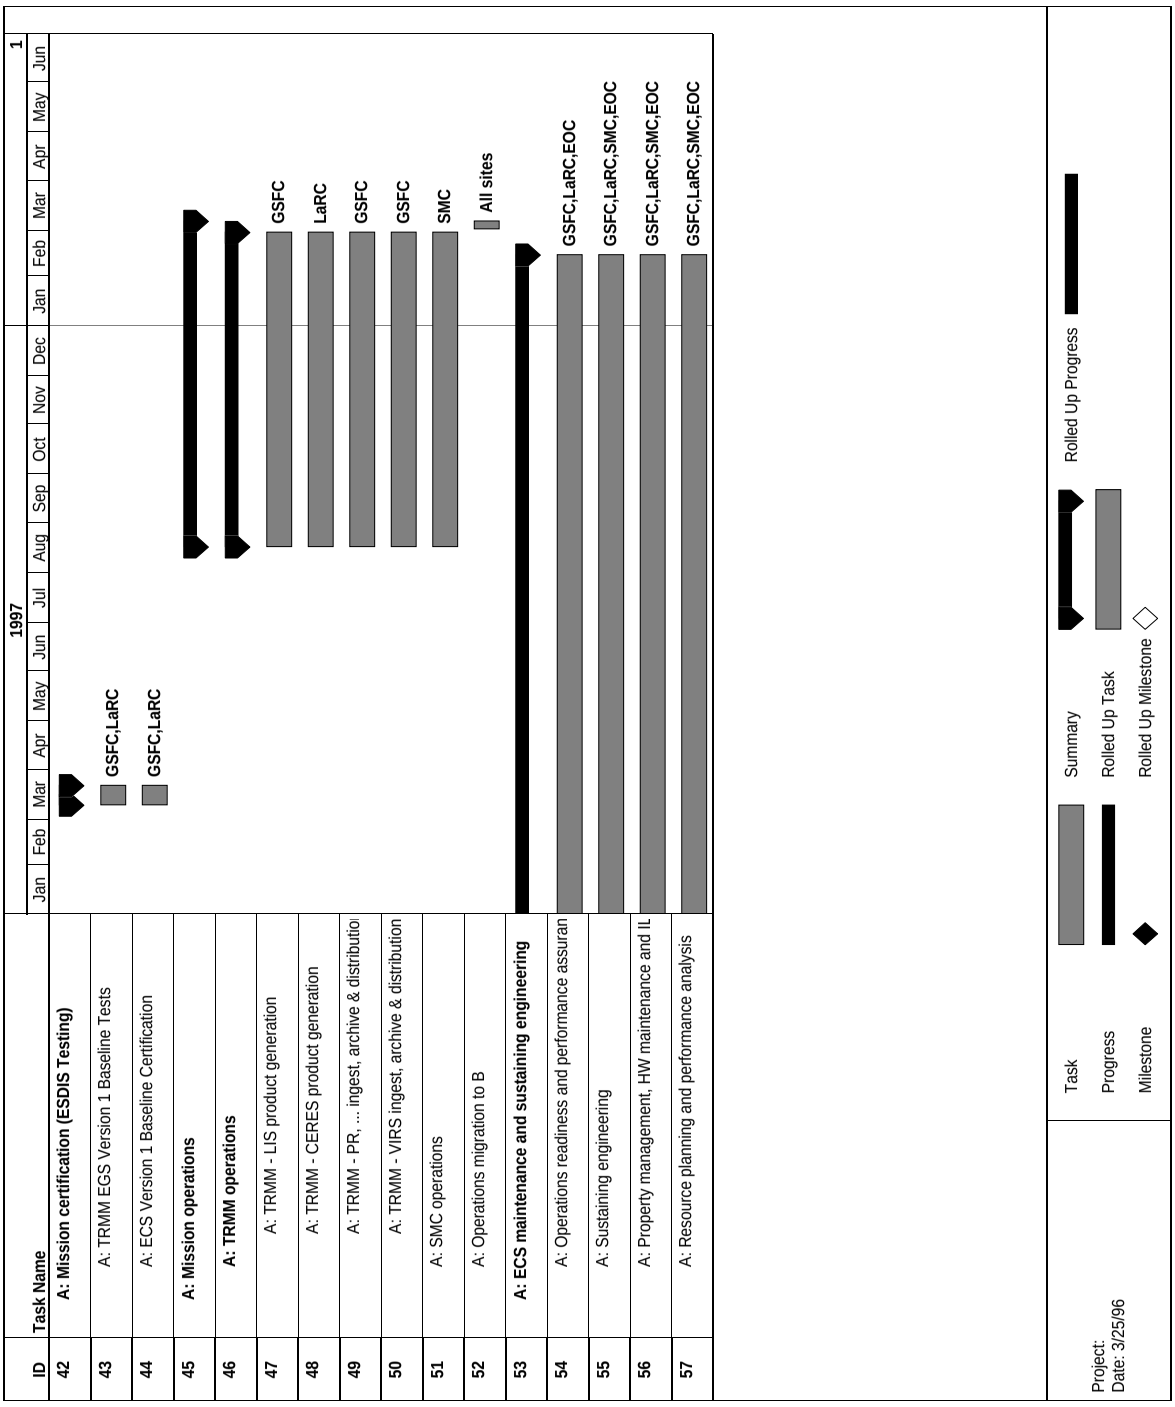

Figure 3-1. Release A Activities During Transiton to B (2 of 2)

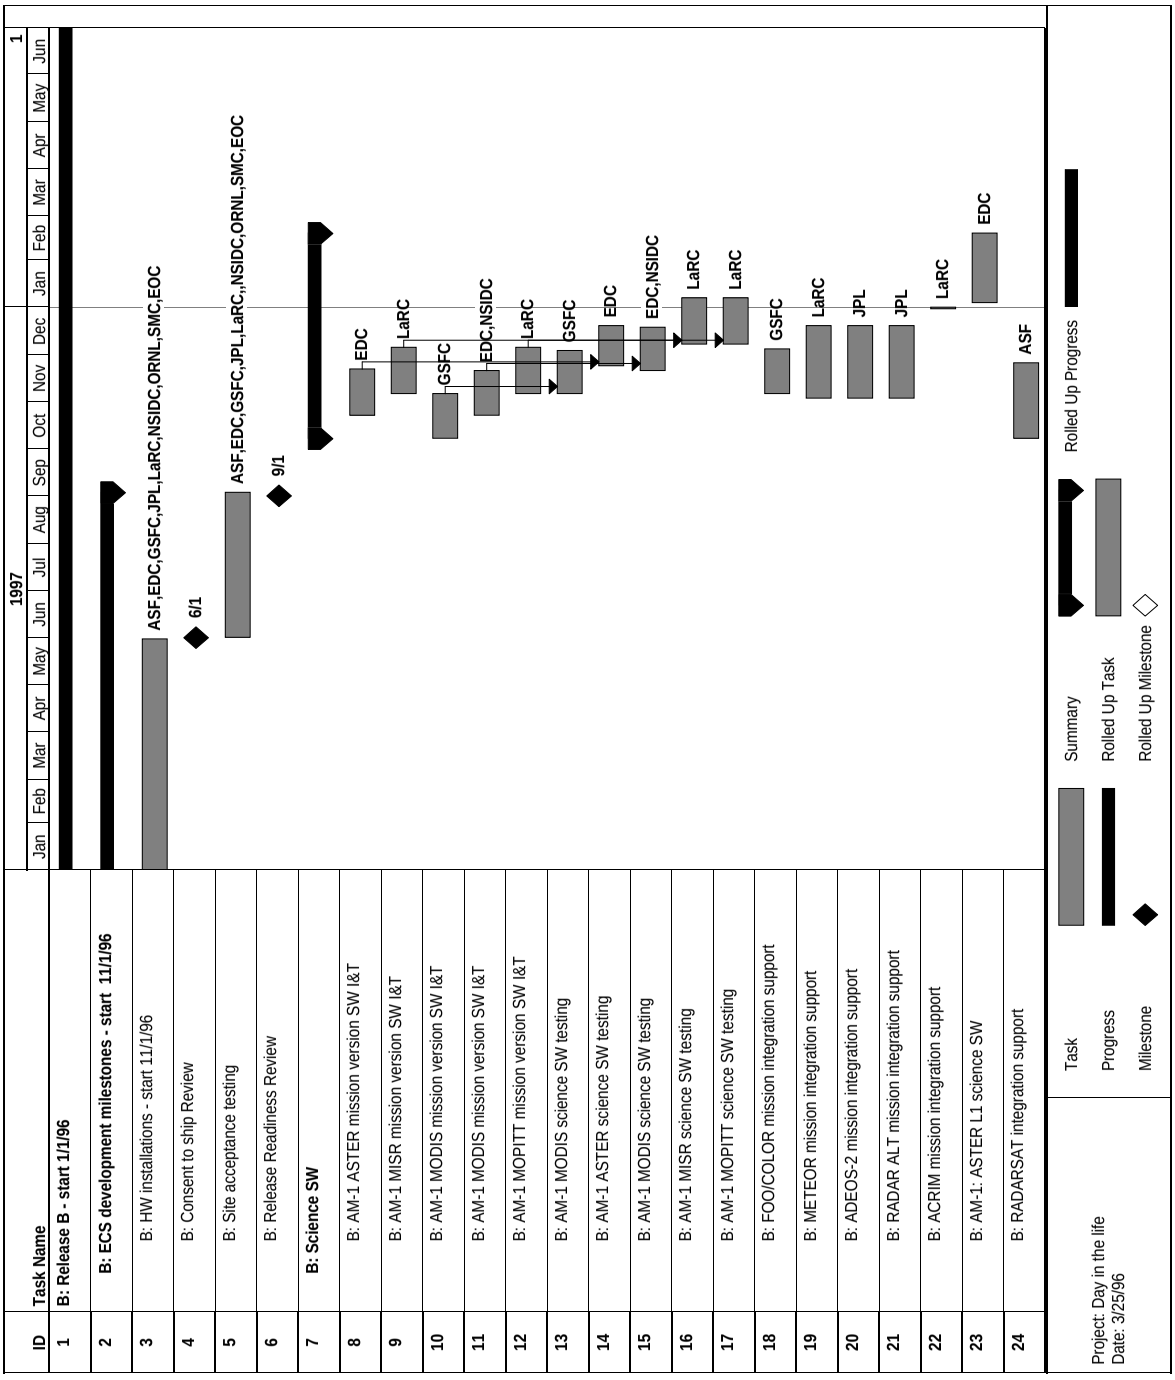

Figure 3-2. Release B Activities During Transition to B (1 of 3)

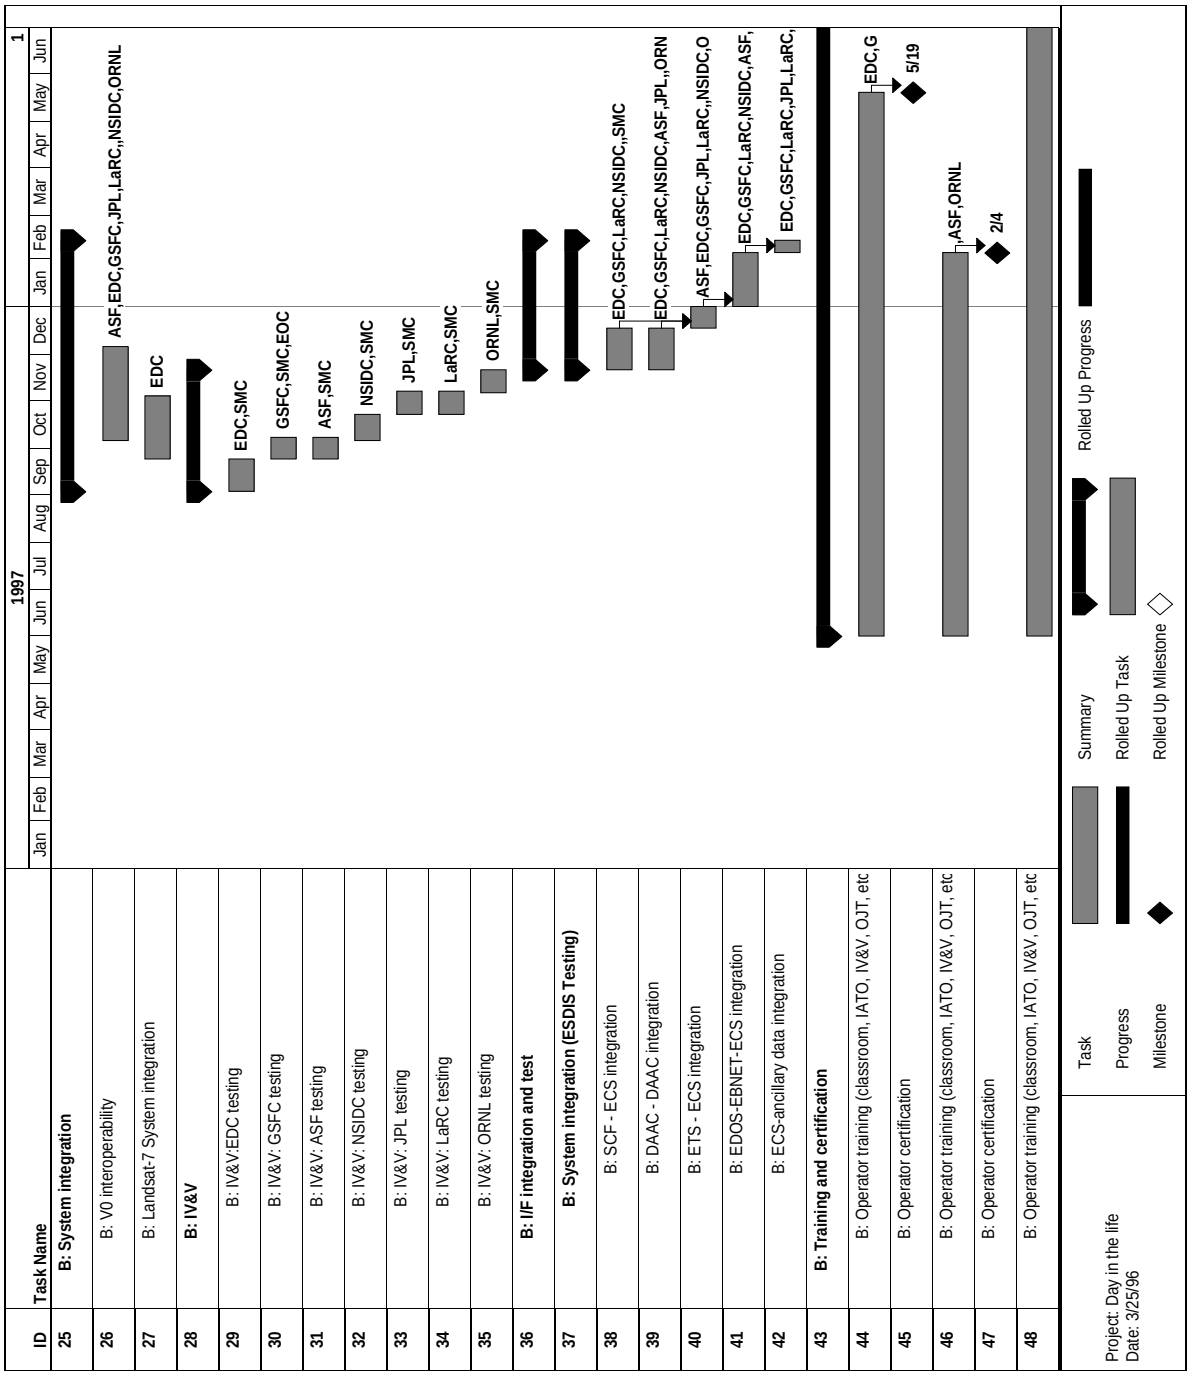

Figure 3-2. Release B Activities During Transition to B (2 of 3)

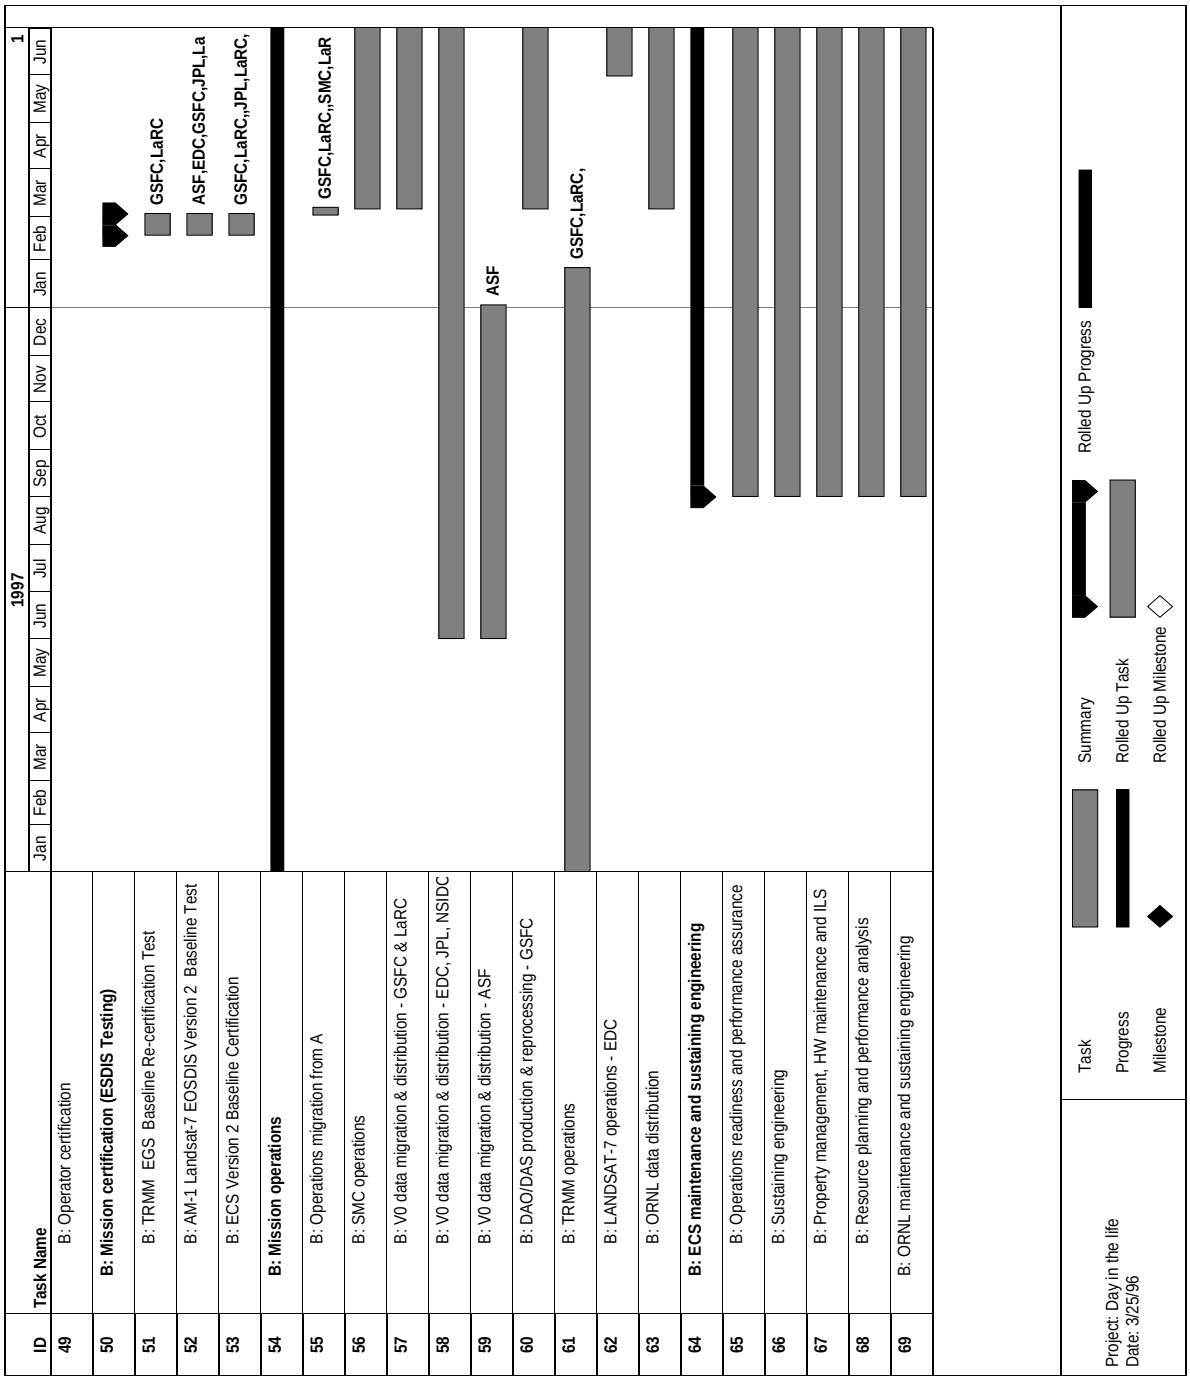

Figure 3-2. Release B Activities During Transition to B (3 of 3)

3.1 Science Software Integration And Test (SSI&T)

3.1.1 Release A/B

TRMM CERES & LIS mission version science software integration and test for both GSFC and Langley are scheduled for one month for Release A starting at the end of January 1997. This is then followed by mission version science software testing at both sites from mid-April 97 until mid-August 1997.

AM-1 ASTER, MISR, MODIS, and MOPITT engineering version of software integration and tests for EDC, LaRC, NSIDC and GSFC are scheduled to run for 3 weeks at each of the named sites. All but ASTER are scheduled to begin in January 1997, with ASTER testing expected during February 1997.

3.1.2 Release B Only Missions

Release B supports the following missions: FOO/COLOR, METEOR, ADEOS-2, RADAR ALT, ACRIM and RADARSAT. The Release B schedule shows science software integration and testing starting as early as October 1997 for some of the missions, and running until mid-February 1998.

3.2 System Integration

3.2.1 VO Interoperability

This topic will be discussed in the next submittal of this white paper.

3.2.2 Landsat 7

This topic will be discussed in the next submittal of this white paper.

3.2.3 Independent Acceptance Test Organization (IATO) Testing Of Release B

Release B provides capabilities to ten ECS sites. These sites are the EOC, SMC and the ECS GSFC DAAC located in Greenbelt, Maryland; the ECS LaRC DAAC in Hampton, Virginia; the ECS EDC DAAC in Sioux Falls, South Dakota; the ECS NSIDC DAAC in Boulder, Colorado; the ECS JPL DAAC in Pasadena, California; the ECS ORNL DAAC in Oak Ridge, Tennessee; the ECS ASF DAAC in Fairbanks, Alaska; and the Consortium for International Earth Science Information Network's (CIESIN) Socio-Economic Data and Applications Center (SEDAC), in Saginaw, Michigan. All DAAC sites except the SEDAC site participate in the Release B acceptance test.

The acceptance testing of Release B capabilities is divided into five major scenario groups: System Management, Push, Pull, Flight Operations, and End-to-End. These scenario groups identify-high level ECS functionality from a user's and operator's viewpoint. Each group is further sub-divided into scenarios that emulate the operations and user environment. Scenarios are further broken down into more manageable test sequences in which test cases that trace to Level-3 requirements are executed.

Detailed traceability of individual scenarios, sequences, and test cases to Level-3 requirements are provided in the Verification Specification (DID 403/VE1, October 1995).

Our current baseline plan calls for conducting Release B acceptance testing during the three month period following the Release B CSR. The plan calls for conducting acceptance tests in four sessions. The first session occurs during the first three week period following CSR at SMC, EOC, GSFC, and LaRC. The second session occurs the following three weeks at EDC and ORNL. The third session occurs the following three weeks at NSIDC, JPL and ASF, with SMC and EOC remaining involved throughout the second and third sessions to participate where mutual testing is required. The final session occurs during the remaining three week period. During the final three weeks an All-Up End-to-End session occurs with all sites participating.

Note that the site order for acceptance testing may change as usage contentions at the sites are examined (refer to section 7.1.3).

The final detailed schedule is to be coordinated with each DAAC site, to minimize disruption to ongoing operations at that DAAC. The final schedule will include dates, times and durations for all formal acceptance testing that occurs at each DAAC site.

3.2.4 Independent Verification And Validation (IV&V) Of Release B

The test plan for the Component Acceptance Test (AT) verifies the Level 3 (L3) requirements assigned to a given release. Component AT precedes the ESDIS Integration and Test (I&T) Program as defined in the Earth Observing System (EOS) Ground System (EGS) Integration, Test & Validation Plan (EITVP). Specific goals for the test plan are to:

Present an overview of the capabilities being tested, including identification of applicable configuration items

Provide an overall schedule of test activities.

Provide descriptions of tests to be executed during EGS testing.

Provide high level test sequences for the applicable sites.

Allocate tests to test sites and provide high level test execution sequences.

Provide additional information regarding test data, personnel, tools, and other related topics.

Testing is scheduled to begin in September 1997 and run until approximately the end of November 1997. The following sites are expected to participate: EDC, GSFC, ASF, NSIDC, JPL, LaRC and ORNL.

3.2.5 Release B Interface Testing

As shown on the Release B M&O Schedule interface testing is scheduled to start mid-November 1997 and run for approximately two months. Five different types of interface testing are shown on the schedule; SCF-ECS, DAAC-DAAC, ETS-ECS, EDOS-EBNET-ECS and ECS-ancillary data locations.

3.3 Release A TRMM Mission

This topic will be discussed in the next submittal of this white paper.

3.4 Operator Training and Certification

For Release B the operator training is a continuation of the Release A training and certification. It will continue for the duration of Release B. Operator training must be completed one and half months before the launches of Landsat 7 and AM-1. This item is noted specifically on both the Release A and the Release B M&O schedules.

3.5 Mission Certification

Mission certification, also known as ESDIS testing, will occur in both Release A as well as Release B. For Release B TRMM EGS Baseline will require re-certification testing. Though the TRMM operations system baseline was established with Release A testing, any changes in the system requires a re-certification of the system's operation mission baseline. This is scheduled for early January 1996. Baseline certification testing will be done from mid-February 1998 until early March 1998. Any further changes to the operational system after successfully completing the Version 2 Baseline Certification (3/98) will involve extensive re-testing.

3.6 Mission Operations

3.6.1 Data Assimilation Office (DAO)

The DAO, located at GSFC and part of the GSFC DAAC (GDAAC), is responsible for developing advanced assimilation algorithms used to produce research-quality assimilated data products. Data from NOAA and other sources are provided to the Data Assimilation System (DAS), at the DAO, in an operational mode. The ECS and the DAO work in coordination with one another to supply ECS with NMC ancillary data required for ECS product generation. The DAO is unique in that it will process data like a DAAC but also develop algorithms like an SCF.

The DAO is not part of Release A, but will be installed, integrated, tested, and made operational after Release B. The nominal processing requirement presently in the technical baseline at Release B is 50 Gflops which translates to on the order of ten workstations. This hardware is then connected to the PDPS LAN with the rest of the science processors at the Goddard DAAC. The need for local mass storage for R&D has also been identified, which is a unique feature in DAAC design. Requiring the use of the GDAAC data server for all R&D work is undesirable since much of the output will be of a tentative nature and will be erased quickly. Local R&D archive may be provided as a dedicated data server, or as mass storage connected to the DAO science processors directly, or as mass storage connected to the rest of the Release B system. This design trade is TBD.

3.6.2 V0 Data Migration

The value of previously archived remote sensing data will be sustained by migrating them into ECS to provide for their continued maintenance. Migration of data from V0 to ECS will begin at

GSFC and LaRC at Release A and continue through transition into Release B. Data at EDC, ASF, JPL, and NSIDC will begin migration at Release B. At ORNL, the existing archives will be retained and only the metadata will be migrated at Release B. The V0 data at MSFC will probably transition to other DAACs starting in the Release A time frame and then migrate into ECS.

Data migration occurs in three steps: data migration engineering, pre-processing, and ECS population. Data migration engineering which includes analysis of data, development of software and/or procedures and the migration of sample data, occurs at the EDF. Pre-processing prepares the data for ECS ingest and will use ECS hardware to perform this function. The actual ingest into ECS is part of normal operations.

Note that the migrations at ASF and ORNL must be completed during the transition time frame, since the data is needed to support operations.

3.6.3 SMC

This topic will be discussed in the next submittal of this white paper.

3.6.4 TRMM

This topic will be discussed in the next submittal of this white paper.

3.6.5 Operations Migration

This topic will be discussed in the next submittal of this white paper.

3.7 EOSDIS Ground System Integration And Test

This topic will be discussed in the next submittal of this white paper.

3.8 ECS Maintenance and Sustaining Engineering Support

3.8.1 Operational Readiness and Performance Assurance

Operational Readiness and Performance Assurance activities with be ongoing throughout this transition period (Mar 97 - Mar 98) in support of both Release A and Release B. This function, at both the system Sustaining Engineering Organization (SEO) and at the DAACs, is the focal point for the planning and monitoring of all training, operational exercises, operator certification, and support for system rehearsals and certification, including the preparation of operations procedures and building of data bases. It is also the focal point for monitoring and reporting on the performance of those aspects of the Release which are currently supporting operations.

Just prior to the transition period, emphasis will be on operational readiness preparations for: V0 data migration; integration and test of the mission versions of the TRMM science software and of the AM-1 version 1 science software; AM-1 I/F testing; and TRMM system certification testing. The activities just mentioned continue into the transiton time frame and are noted here to place them in their proper schedule context.

March - April 97 -- The focus will be on the development of TRMM operations procedures and data bases, on the preparation for Release B and AM-1 training and in support of TRMM system certification testing. At the same time, attention will also be focused on the performance of the Release A operational activities (V0 data migration, SSI&T and AM-1 IF testing).

May - June 97 -- Focus will be on TRMM console training and operational exercises, leading to operator certification; to updating the TRMM procedures and data bases; and the continued support of TRMM system certification testing. During this period Release B procedures will be updated in preparation for Release B and AM-1 classroom and console training, otherwise known as on-the-job (OJT) training, that occurs during the support of the integration and test of Release B, scheduled to begin in June 97. At the same time, attention will continue to be focused on the performance of the Release A operational activities (V0 data migration, SSI&T and AM-1 IF testing).

July 97 -- Support TRMM system rehearsals, leading to TRMM M&O readiness (HW, SW, operations staff and operational procedures and databases. Continue OJT and support of Release B integration and test. Continue focus on the performance of the Release A operational activities (V0 data migration and AM-1 IF testing).

August 97 -- Update Release B procedures and data bases. Plan Release B Site and System exercises. Continue OJT and support of Release B integration and test. Continue focus on the performance of the Release A operational activities (V0 data migration, AM-1 IF testing and TRMM operations, which are scheduled to begin this month).

September - November 97 -- Support Release B IV&V, verify Release B procedures and data bases during operations exercises. Continue focus on the performance of the Release A operational activities (V0 data migration, SSI&T, AM-1 IF testing and TRMM operations.)

November 97 - January 98 -- Support site and system level exercises of Release B. Transition of Release A operations functions and data to the Release B system and parallel operations evaluations will occur. Begin development of AM-1 operational procedures. Continue focus on the performance of the Release A operational activities (V0 data migration, SSI&T, AM-1 IF testing, and TRMM operations.)

January - June 98 -- Execute operations on Release B. Continue focus on the performance of the operational activities (V0 data migration, AM-1 IF testing, and TRMM operations) which are now on Release B. Prepare for AM-1 launch and execution of AM-1 science processing with operator training and operations exercises to verify AM-1 procedures, databases, and system rehearsals to verify system readiness.

3.8.2 Sustaining Engineering

ECS Sustaining Engineering begins in Release A (see Release A schedule) under control of the M&O Organization. Transition supporting software faces several time events of particular note. Ideally, transition supporting software should be installed and tested before March 1997. This is when TRMM mission baseline certification is scheduled for completion. If this date is not met, then schedule impact is likely due to additional EDSIS re-testing and re-certification of the

mission system baseline. Whether additional testing is necessary after March 1997, all testing must be completed no later than June 1997 due to the Release A system freeze.

A second deadline of note that transition software must address is delivery before Release B CSR, scheduled for June 1997. At that time Release B software is expected to be deployed to the sites, making it very difficult to make any further software changes. During this time we are also under going operator certification. It should be noted that there is a Release A system freeze starting June 1997 and no additonal changes will be permitted.

3.8.3 Property Management

This topic will be discussed in the next submittal of this white paper.

4. Release A Operational Configurations

ECS Release A will be deployed at three DAAC locations: Goddard Space Flight Center (GSFC), Langley Research Center (LaRC), and the EROS Data Center (EDC). Release A is also deployed to the System Monitoring And Coordination (SMC) center at GSFC. This section of the document describes the hardware and network components along with their associated COTS software for each Release A ECS DAAC and the SMC. Each section includes a brief description of the DAAC configuration and lists the ECS hardware that it contains. The remaining information is found in the associated hardware diagram for each DAAC as identified within the section write-up.

The next release of this paper will include more detailed software to hardware mappings for each Release A site, showing the custom software executables and the versions of the COTS products that are hosted by each processor. For version 1, the Release A configuration diagrams are shown. The additional hardware and COTS for Release B is also shown, in grey shading.

4.1 GSFC ECS DAAC

The GSFC ECS DAAC is the largest Release A configuration. Figure 4-1, GSFC ECS DAAC Hardware Configuration, illustrates the hardware components and the associated COTS software. The hardware components which compose the GSFC ECS DAAC include SPRHW, AITHW, ICLHW, PLNHW, MSS, CSS, ISS, DMGHW, ACMHW, DIPHW, DDSHW, DRPHW, and WKSHW.

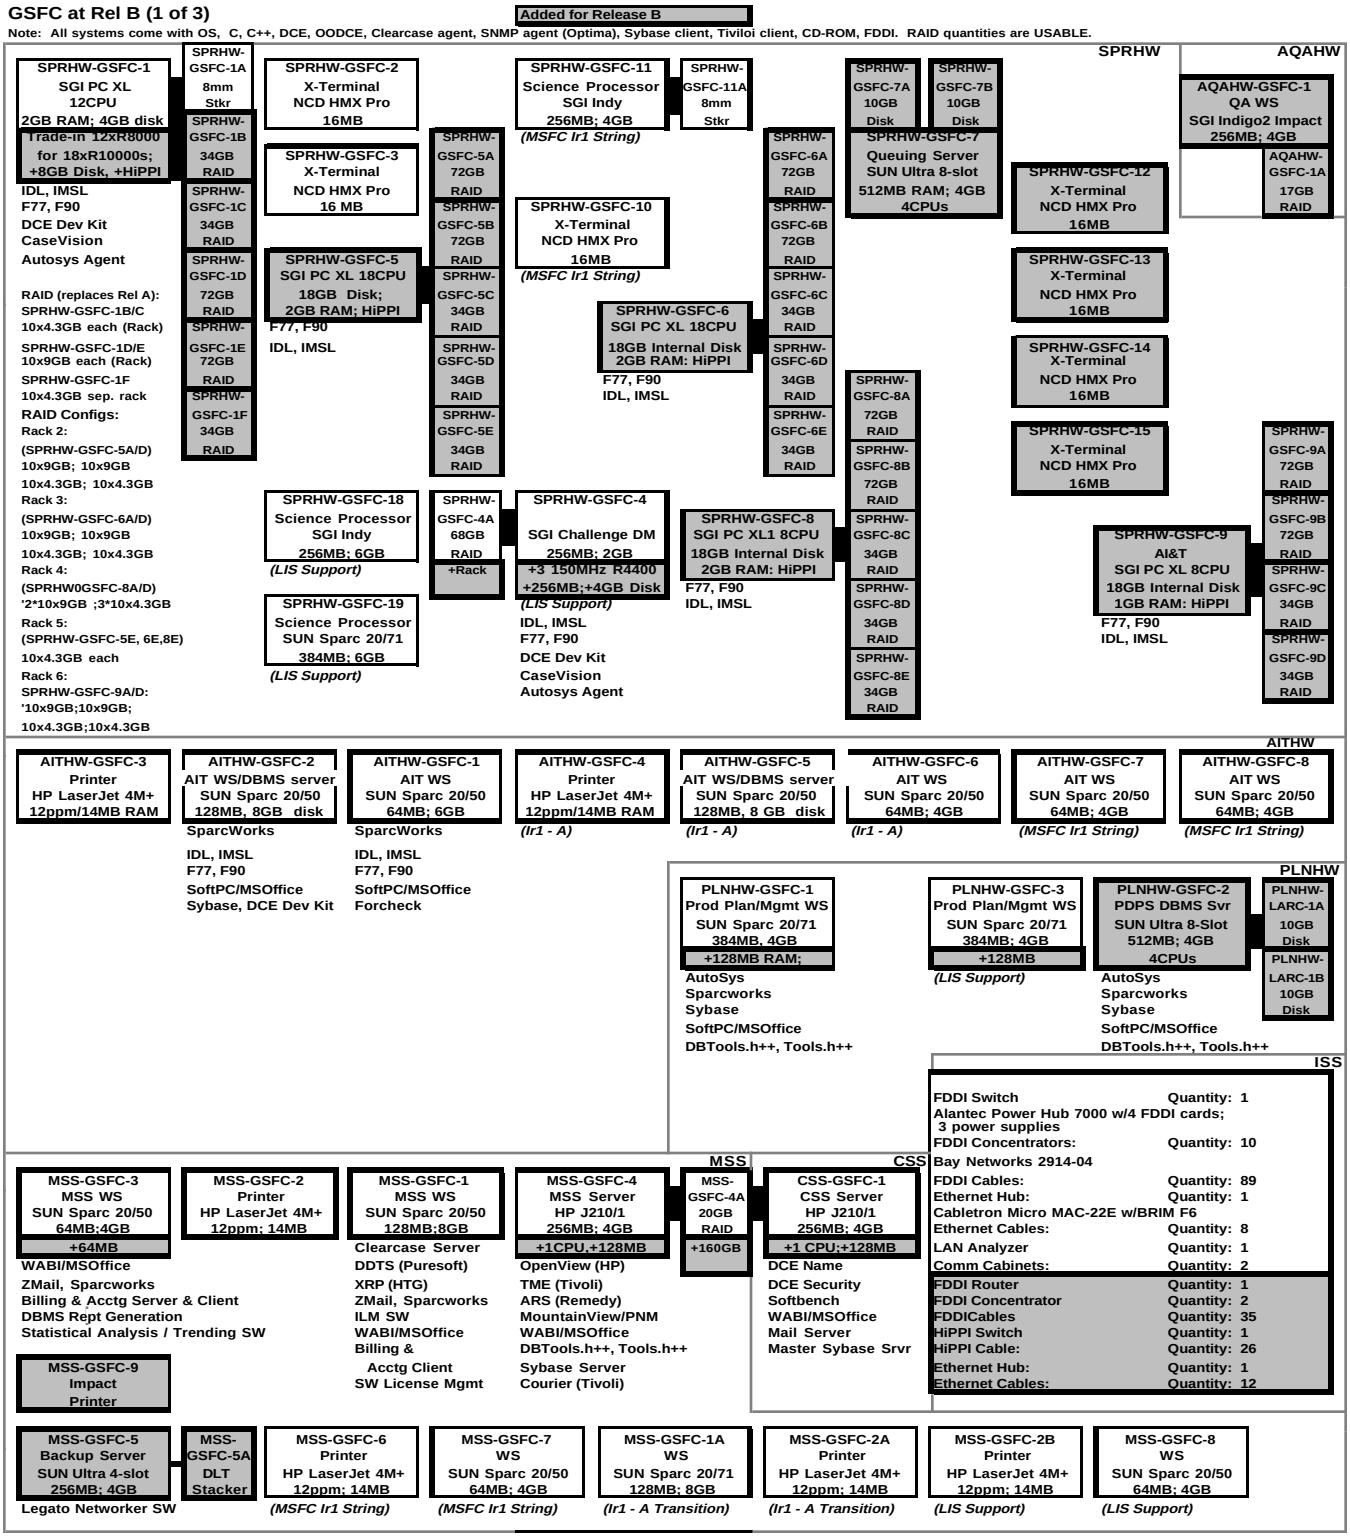

Figure 4-1. GSFC ECS DAAC Hardware Configuration (1 of 3)

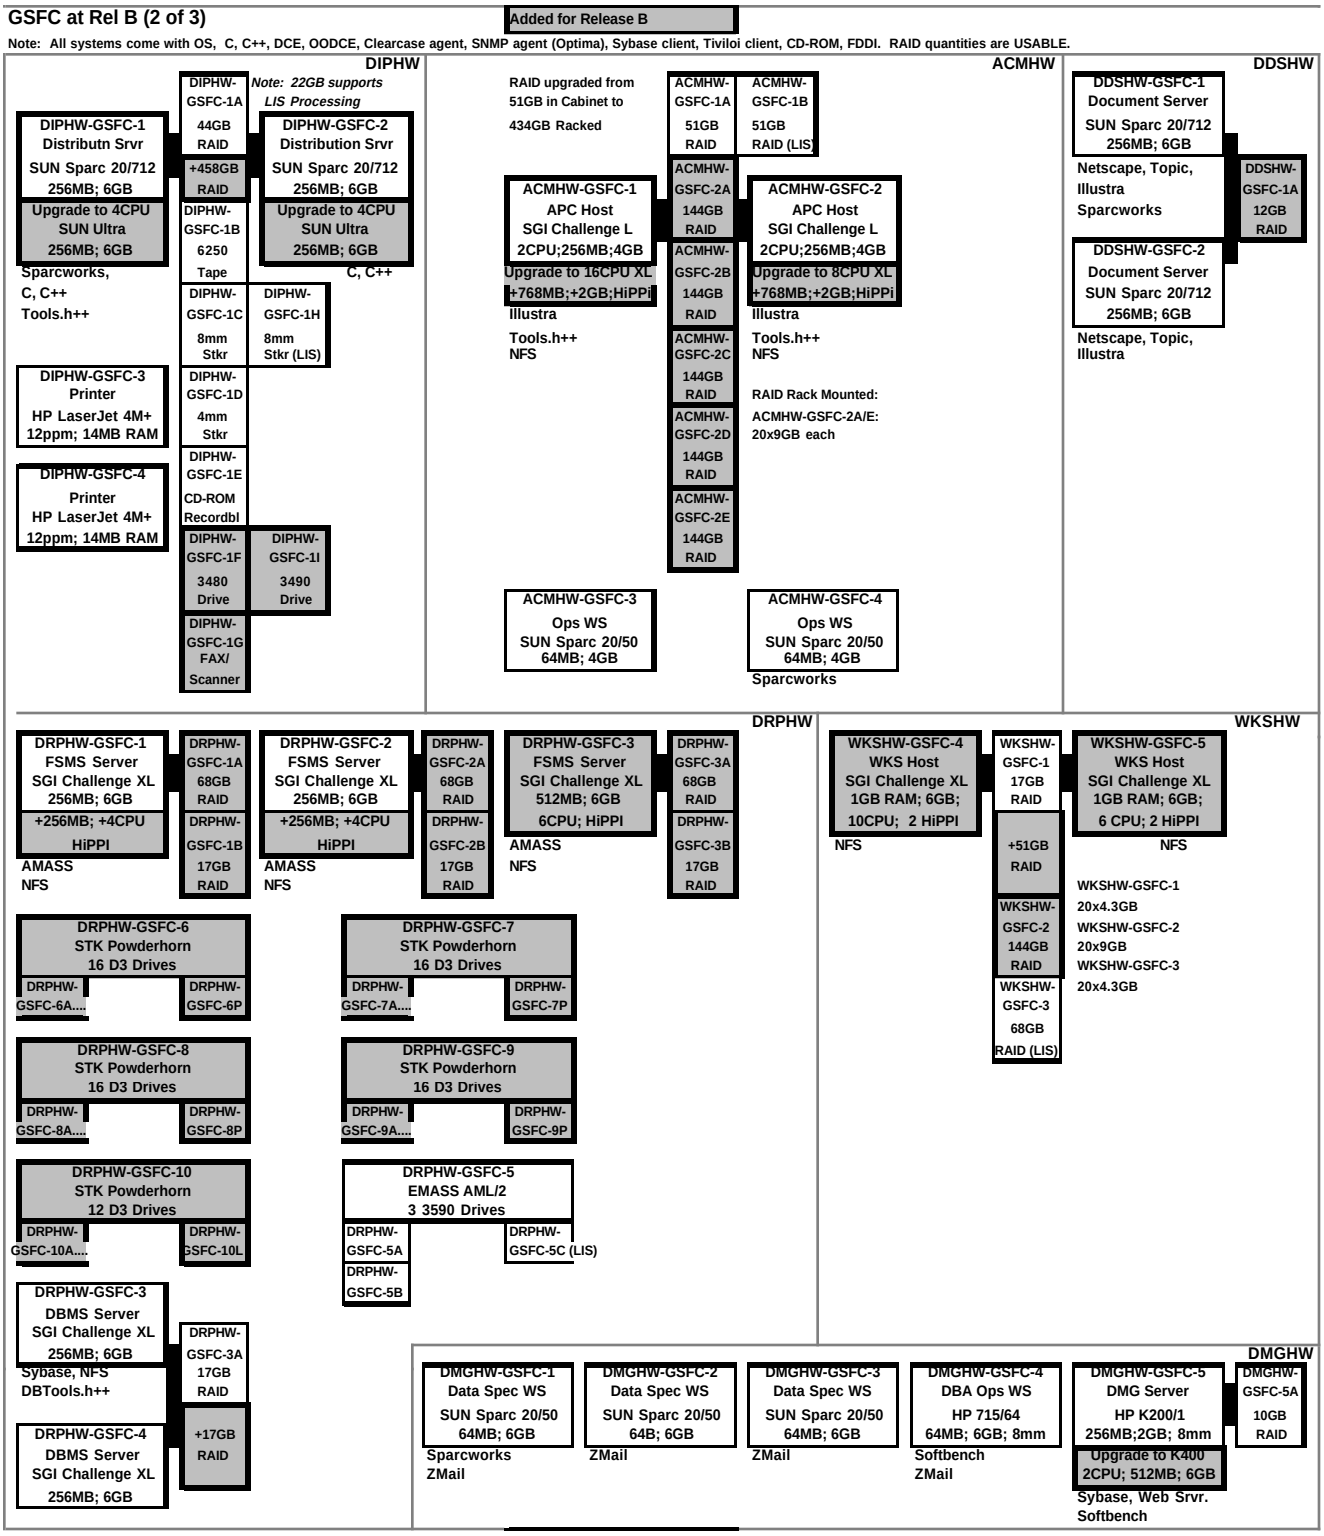

Figure 4-1. GSFC ECS DAAC Hardware Configuration (2 of 3)

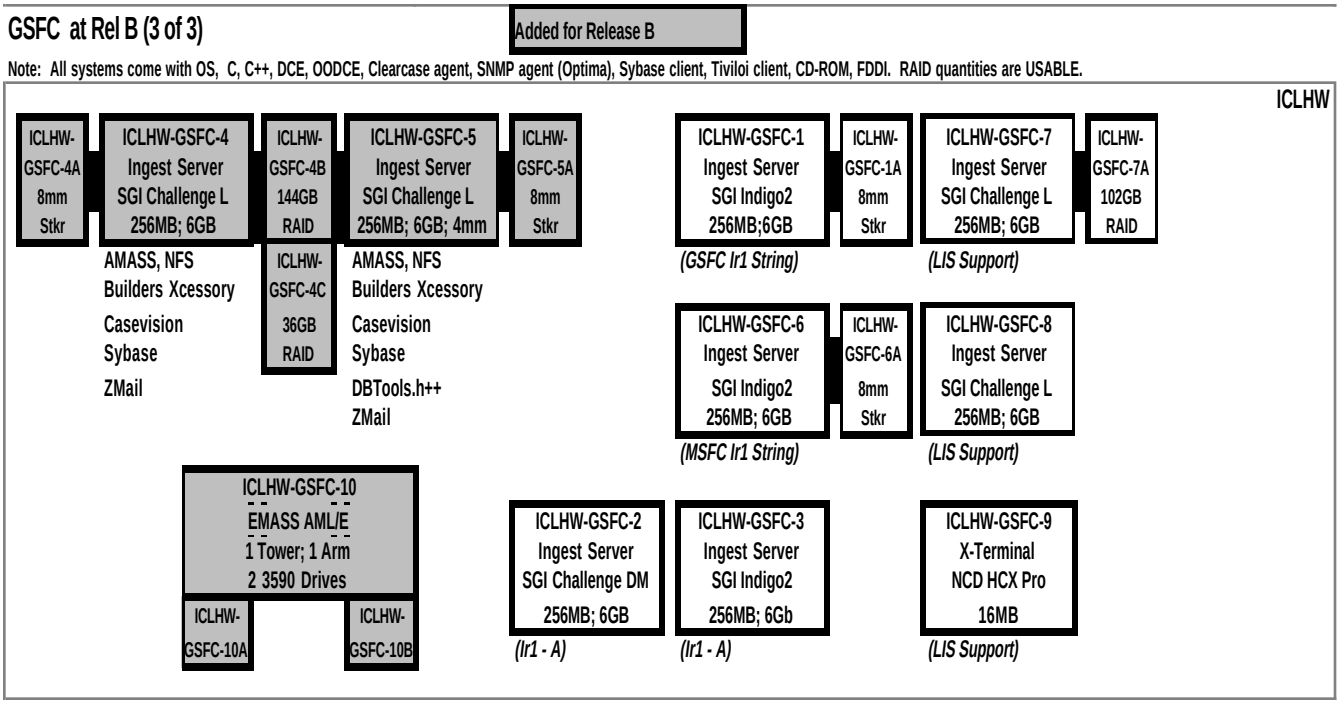

Figure 4-1. GSFC ECS DAAC Hardware Configuration (3 of 3)

4.2 LaRC ECS DAAC

The LaRC ECS DAAC is also a large Release A configuration. As illustrated in Figure 4-2, LaRC ECS DAAC Hardware Configuration, the LaRC ECS DAAC includes hardware and COTS software for the following HWCIs: SPRHW, AITHW, AQAHW, PLNHW, MSS, CSS, ISS, DMGHW, ICLHW, ACMHW, DIPHW, WKSHW, DDSHW, and DRPHW.

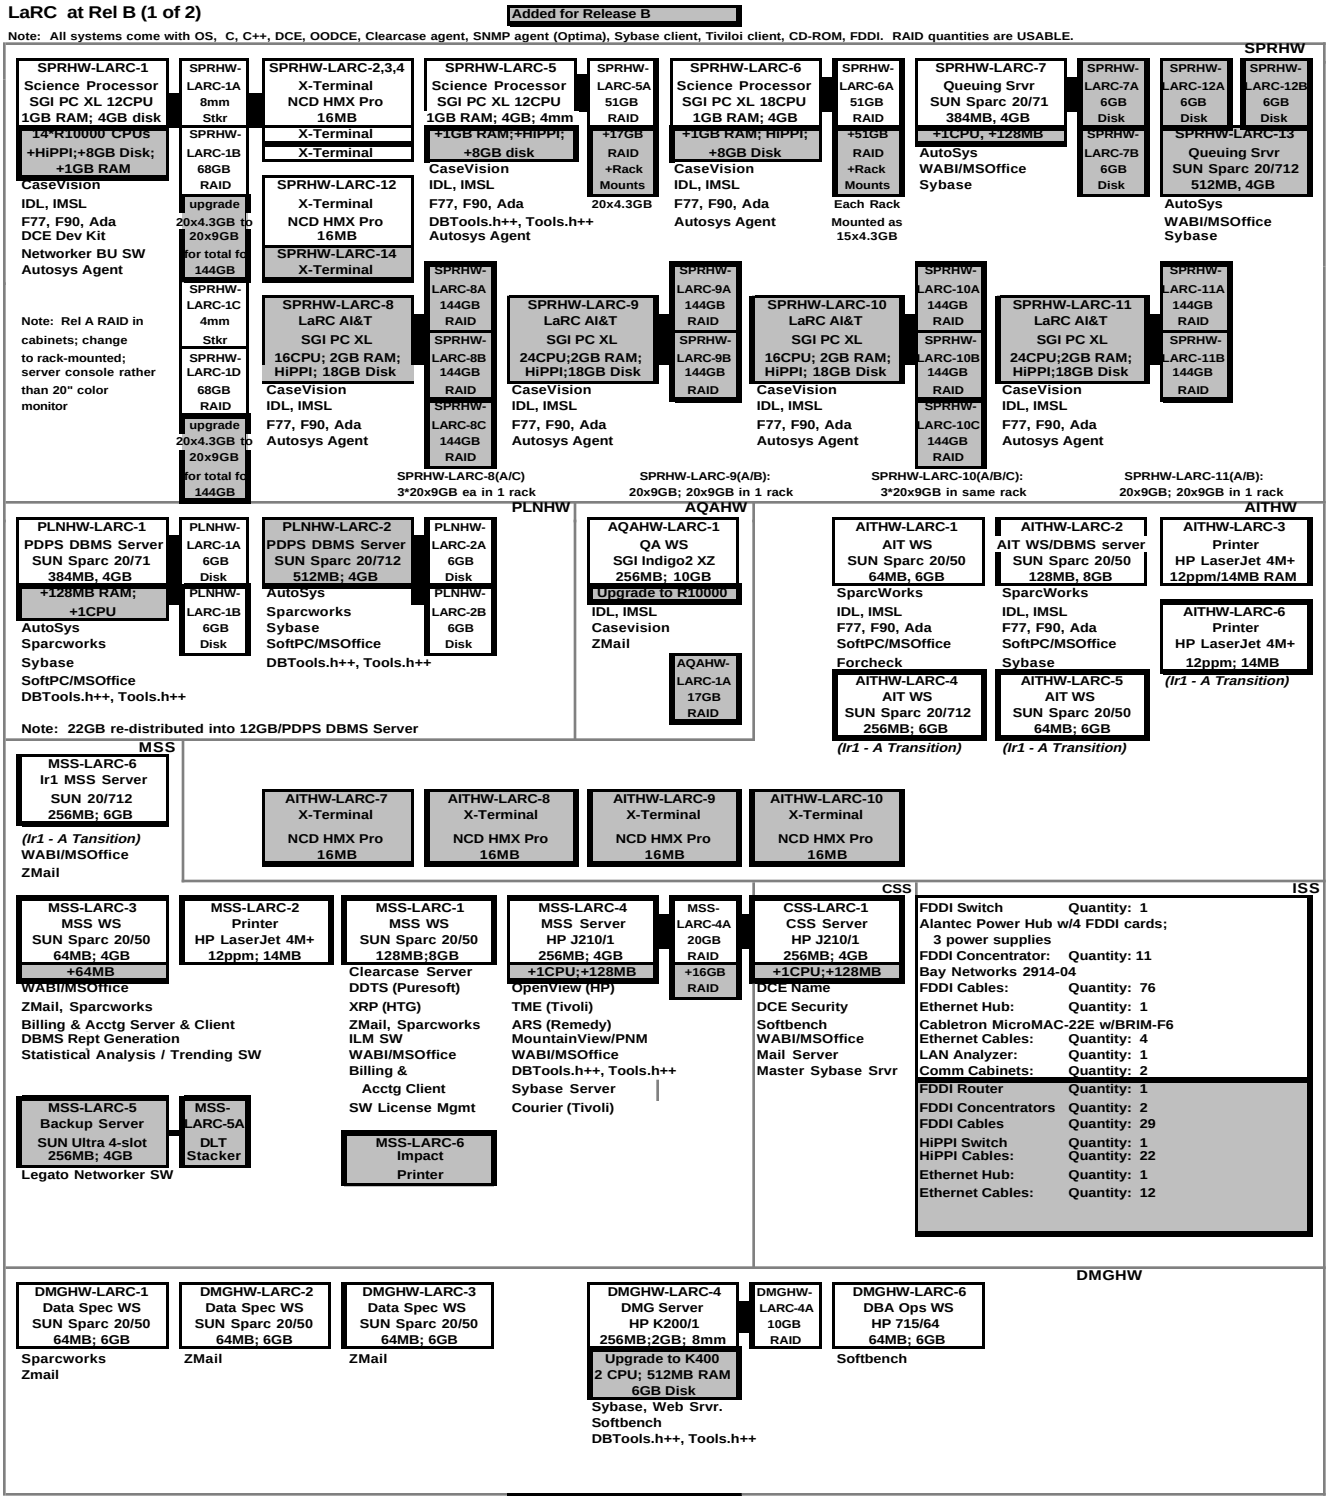

Figure 4-2. LaRC ECS DAAC Hardware Configuration (1 of 2)

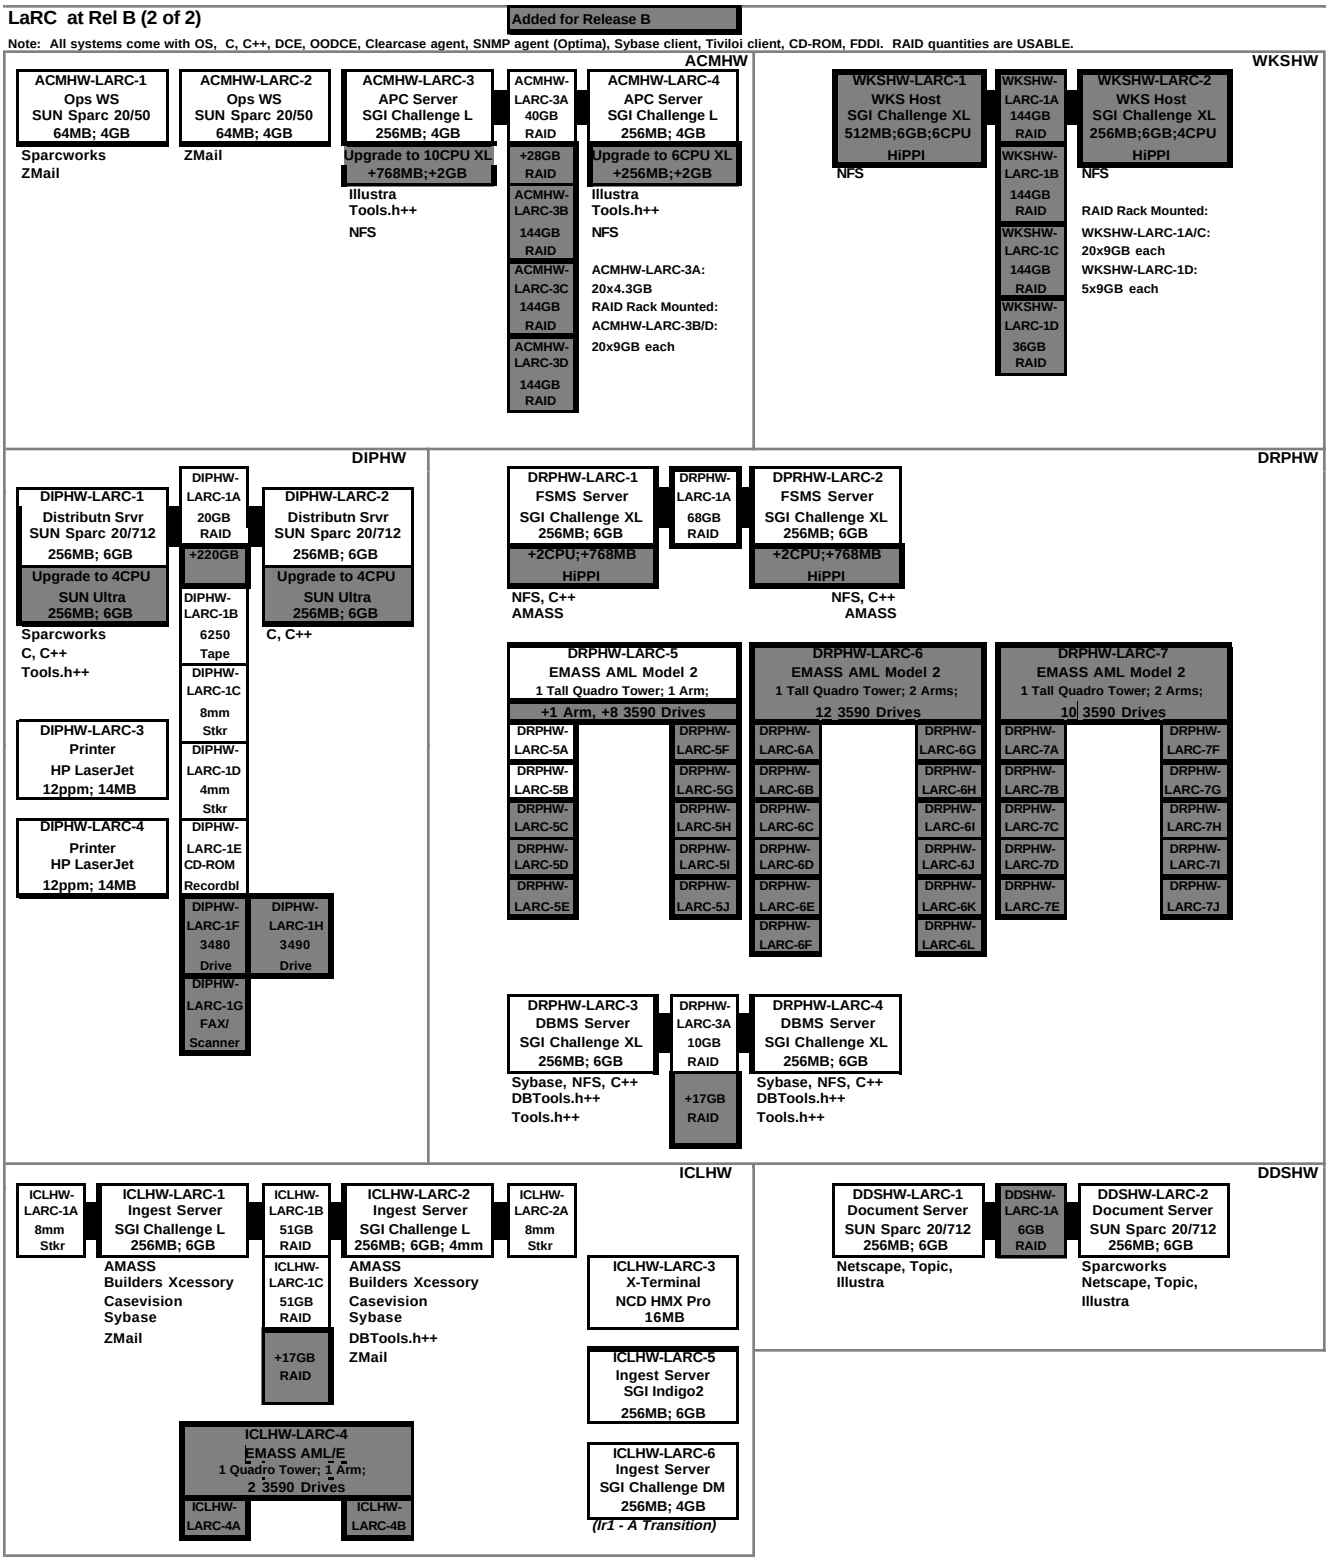

Figure 4-2. LaRC ECS DAAC Hardware Configuration (2 of 2)

4.3 EDC ECS DAAC

The EDC ECS DAAC is the smallest of the Release A configurations. As illustrated in Figure 4∞ 3, EDC ECS DAAC Hardware Configuration, the EDC ECS DAAC includes the SPRHW, AITHW, ICLHW, MSS, and also the ASTER LUT hardware and associated COTS software. Although EDC is not a production site in Release A, it supports interface testing and SSI&T.

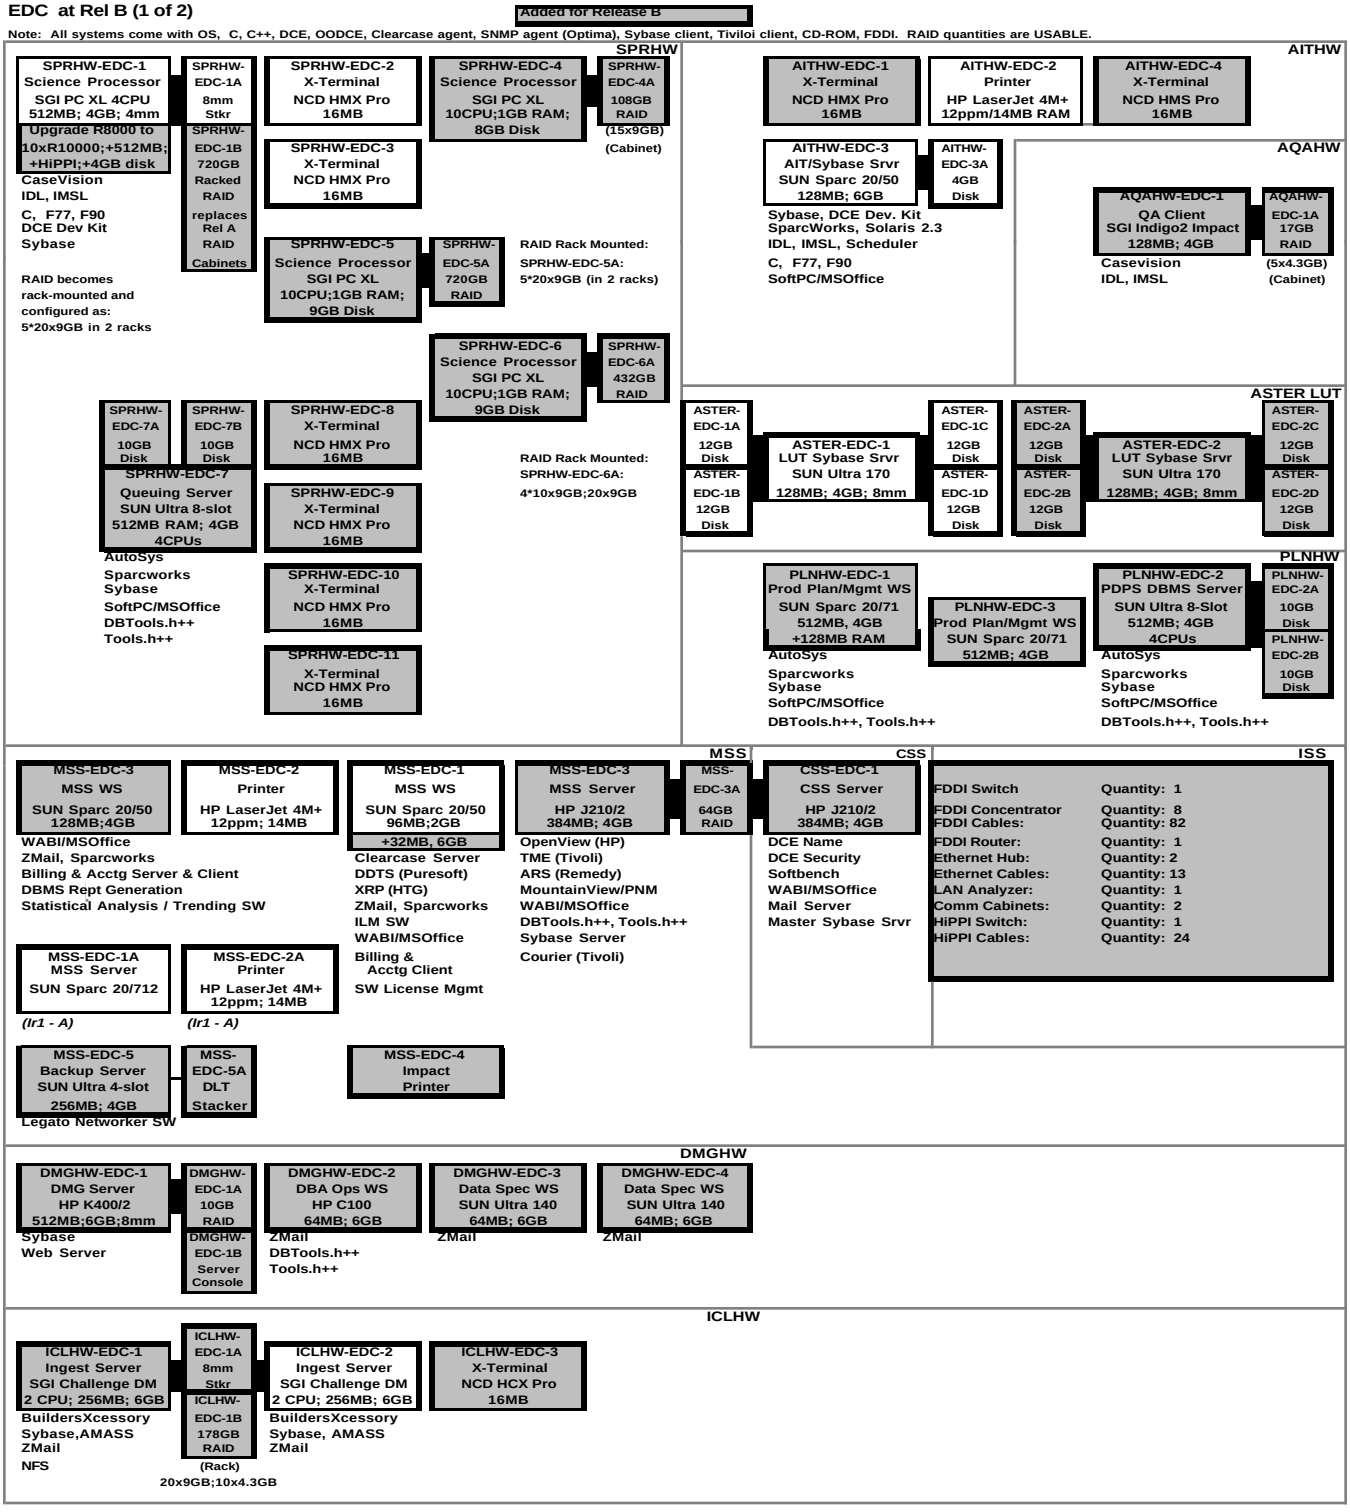

Figure 4-3. EDC ECS DAAC Hardware Configuration (1 of 2)

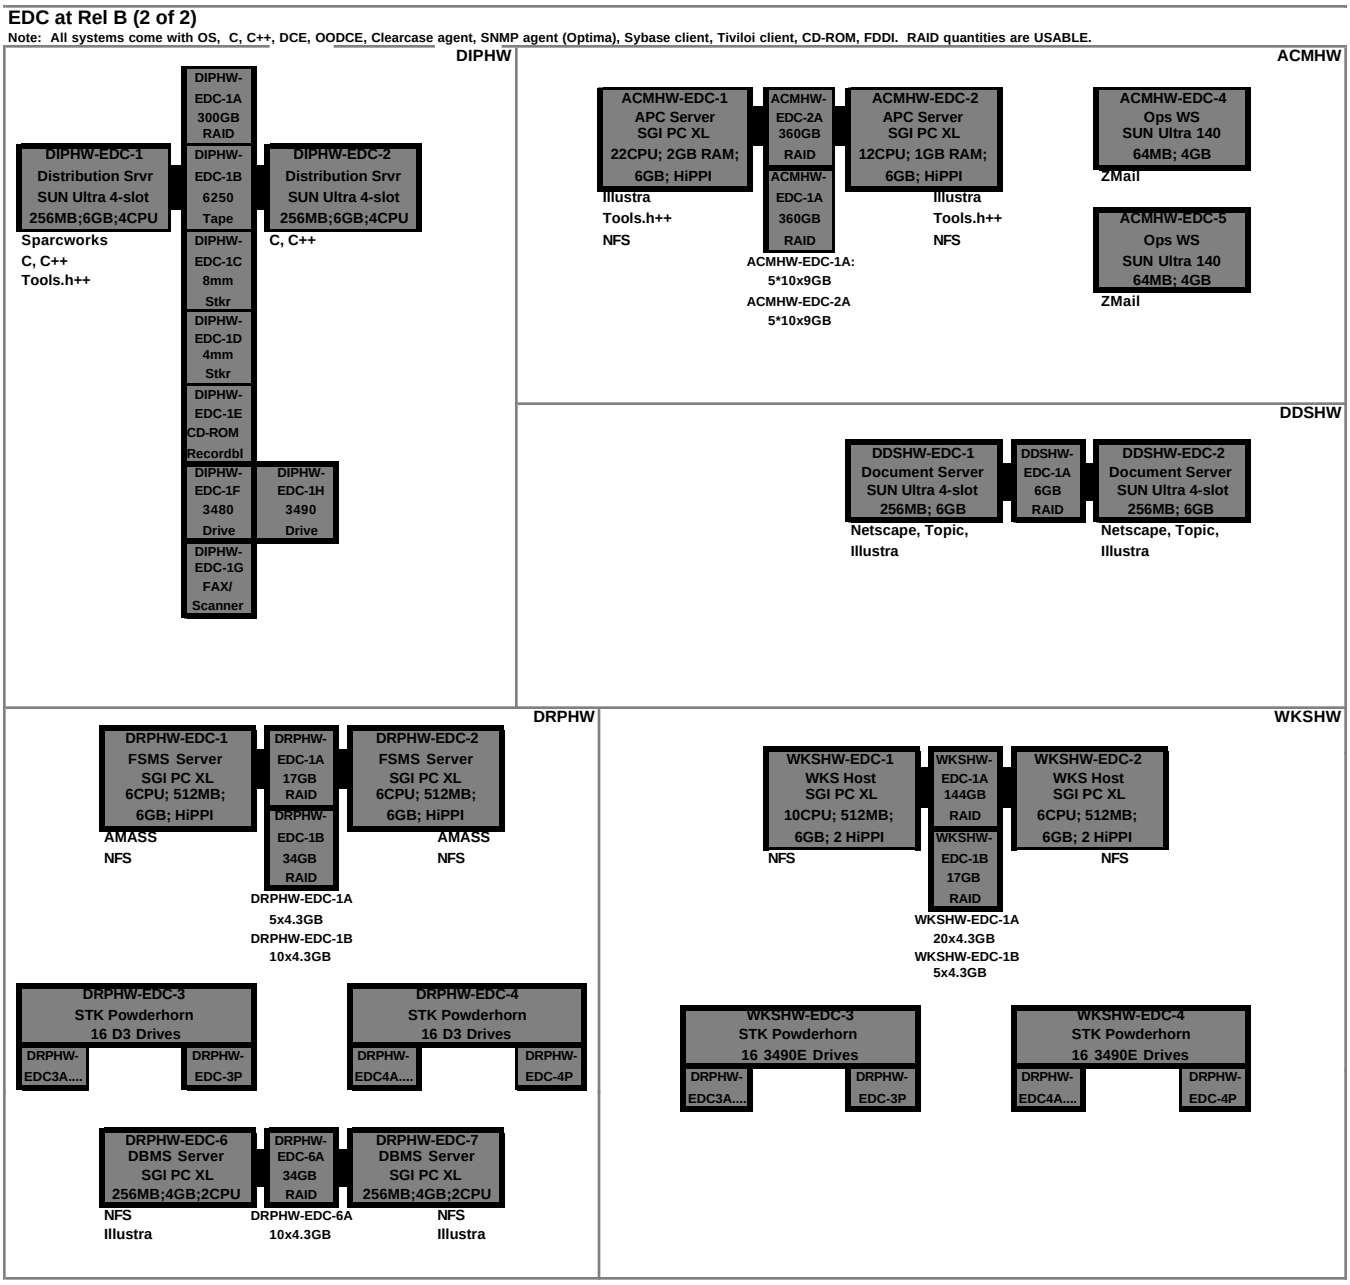

Figure 4-3. EDC ECS DAAC Hardware Configuration (2 of 2)

4.4 ECS SMC

The ECS SMC comprises the MSS, CSS, and ISS subsystems used for system monitoring and coordination, as shown in Figure 4-4 ECS SMC Hardware Configuration.

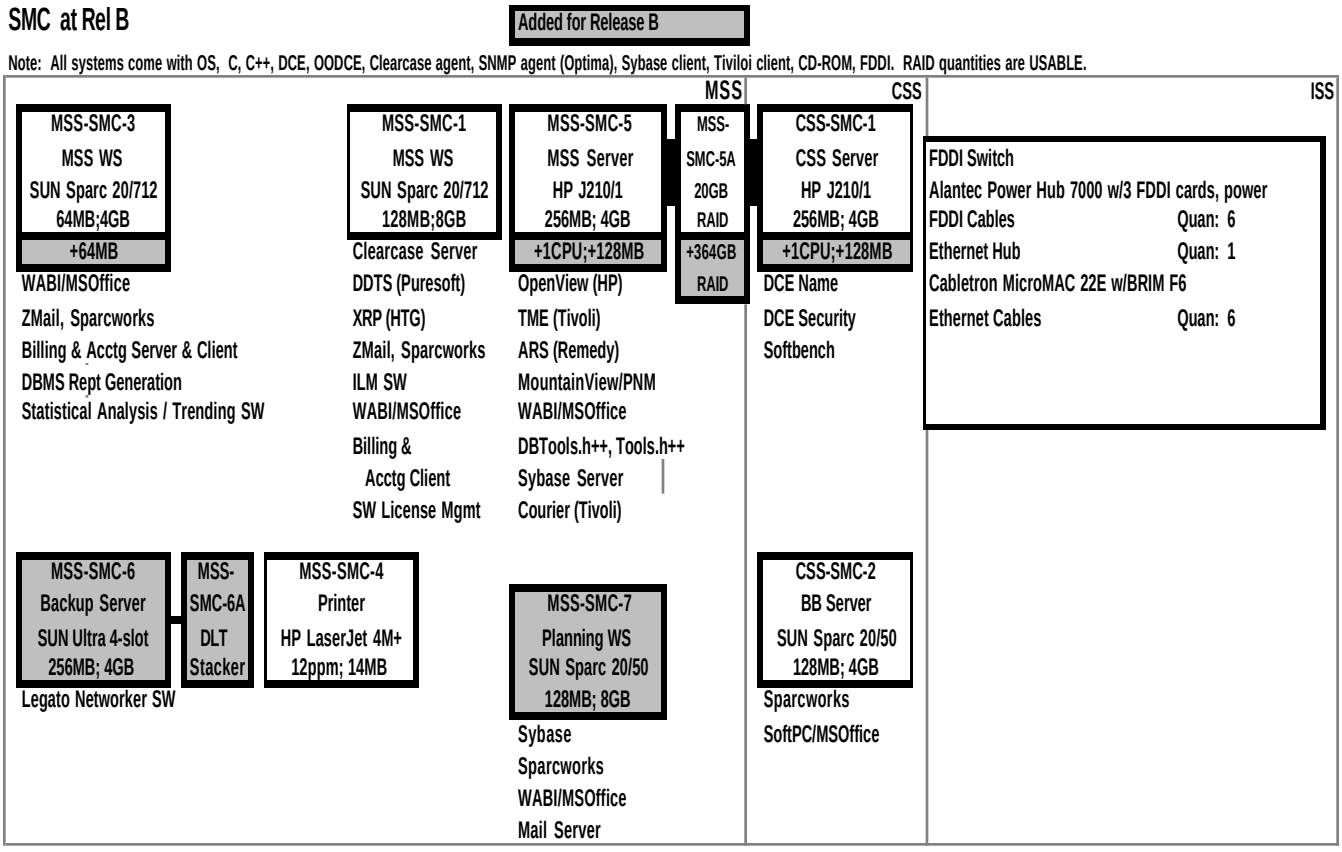

Figure 4-4. ECS SMC Hardware Configuration

5. Release B Operational Configurations

The following sites are added to the ECS system in the Release B timeframe. This section of the document describes the hardware and network components along with their associated COTS software for each new Release B ECS DAAC. Each section includes a brief description of the DAAC configuration and lists the ECS hardware that it contains. The remaining information is found in the associated hardware diagram for each DAAC as identified within the section write∞ up.

5.1 NSIDC ECS DAAC

The NSIDC ECS DAAC is illustrated in Figure 5-1, NSIDC ECS DAAC Hardware Configuration. The NSIDC ECS DAAC includes hardware and COTS software for the following HWCIs: SPRHW, AITHW, AQAHW, PLNHW, MSS, CSS, ISS, DMGHW, ACMHW, DIPHW, WKSHW, DDSHW, and DRPHW. The Ingest software is hosted by the ACMHW CI at NSIDC.

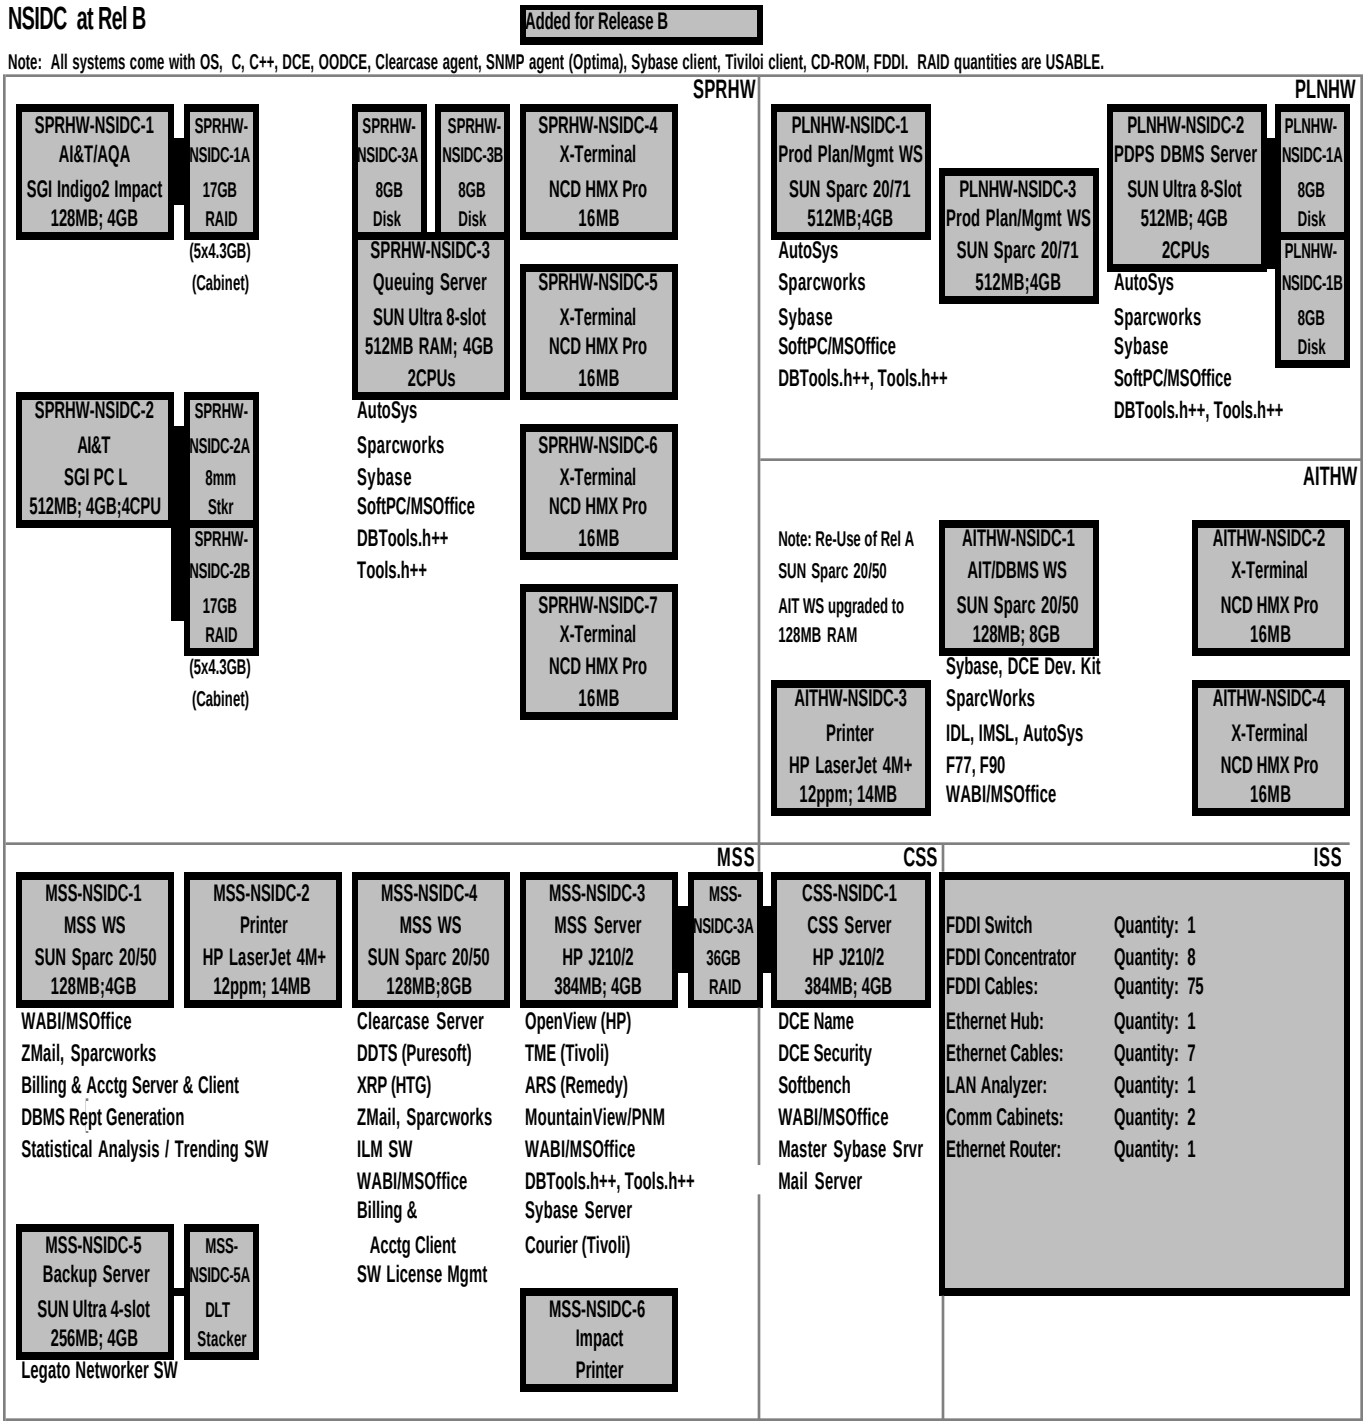

Figure 5-1. NSIDC ECS DAAC Hardware Configuration (1 of 2)

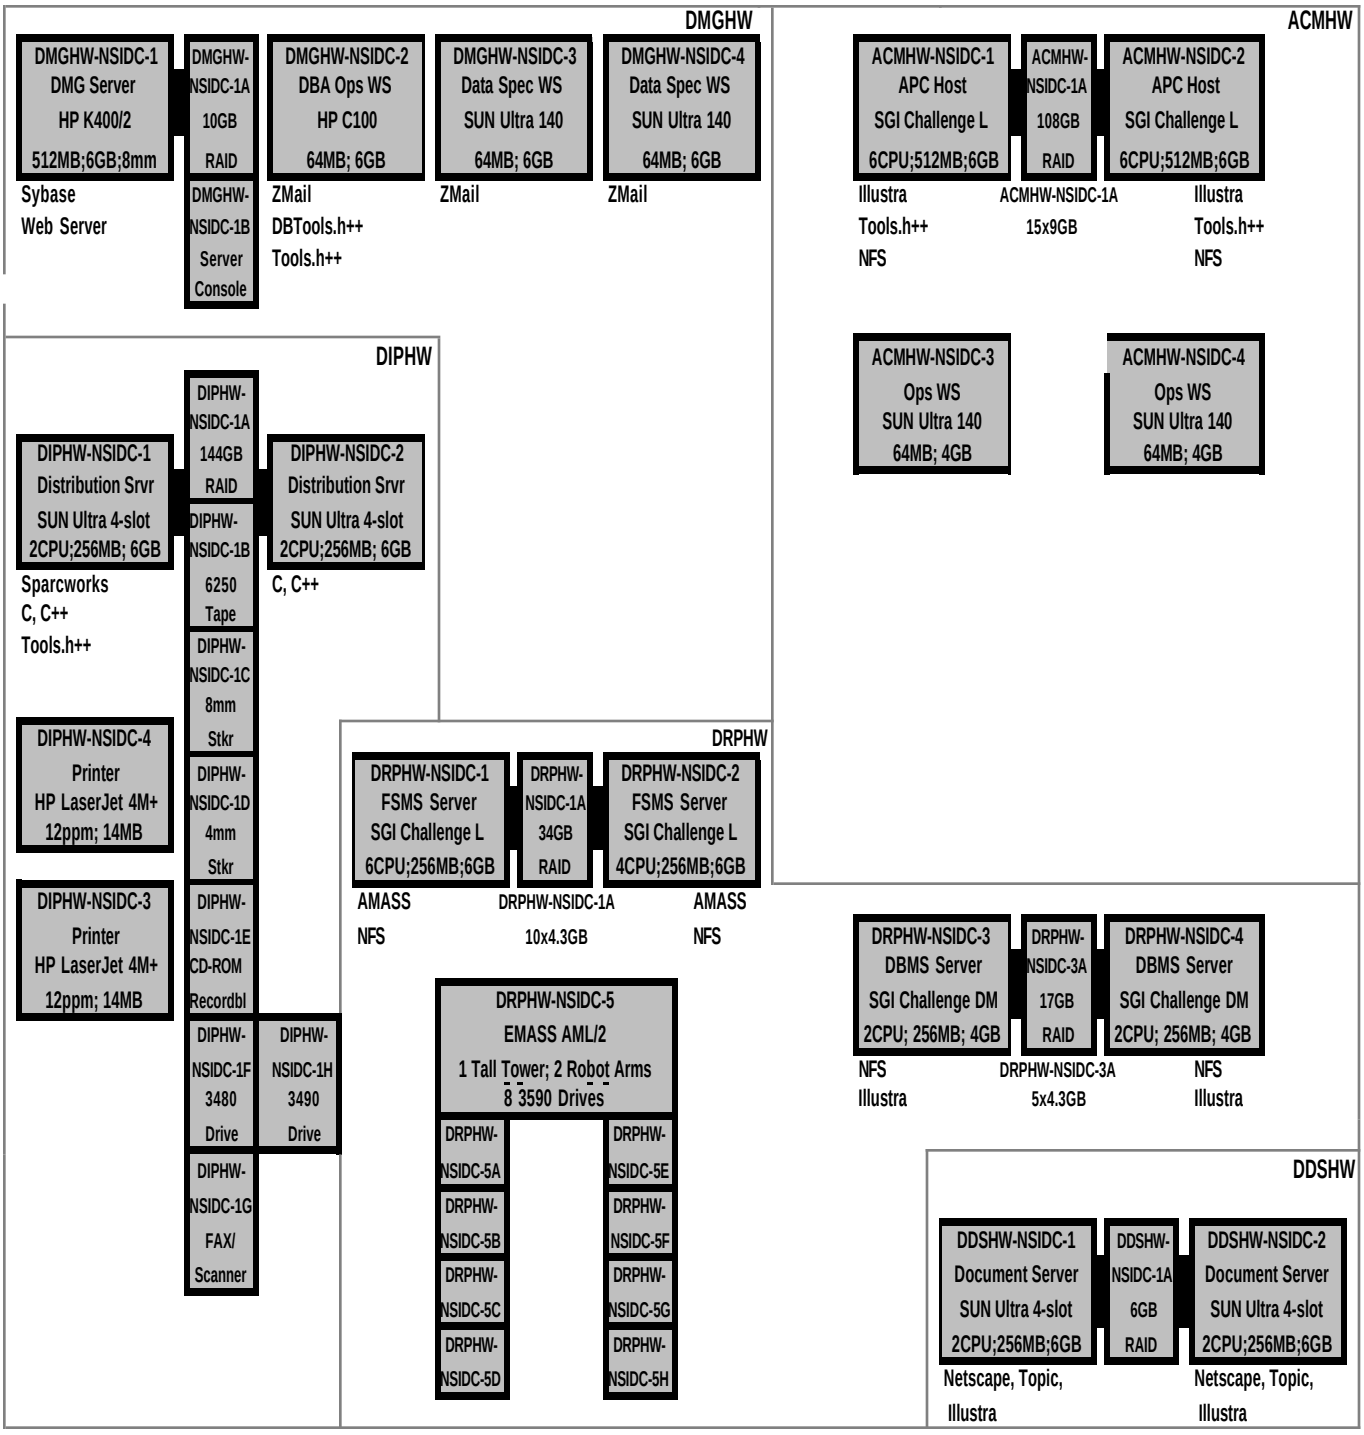

Figure 5-1. NSIDC ECS DAAC Hardware Configuration (2 of 2)

5.2 JPL ECS DAAC

The JPL ECS DAAC is illustrated in Figure 5-2, JPL ECS DAAC Hardware Configuration. The JPL ECS DAAC includes hardware and COTS software for the following HWCIs: SPRHW, AITHW, AQAHW, PLNHW, MSS, CSS, ISS, DMGHW, ACMHW, DIPHW, ICLHW, WKSHW, DDSHW, and DRPHW.

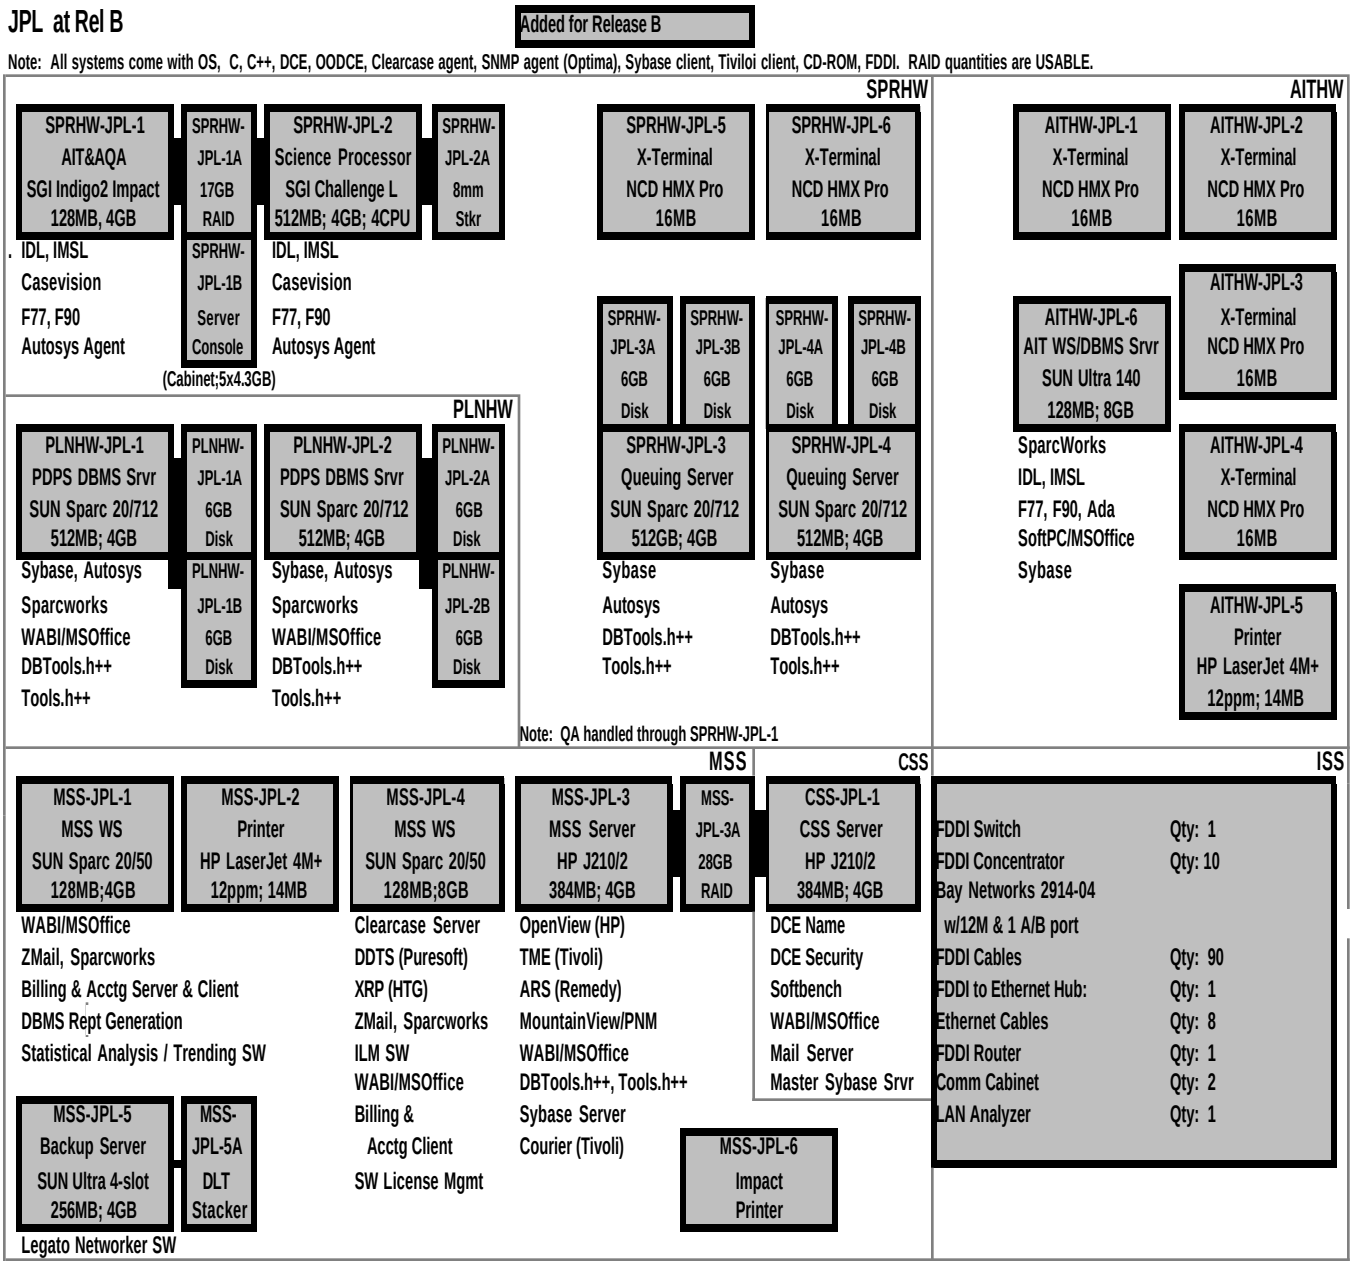

Figure 5-2. JPL ECS DAAC Hardware Configuration (1 of 2)

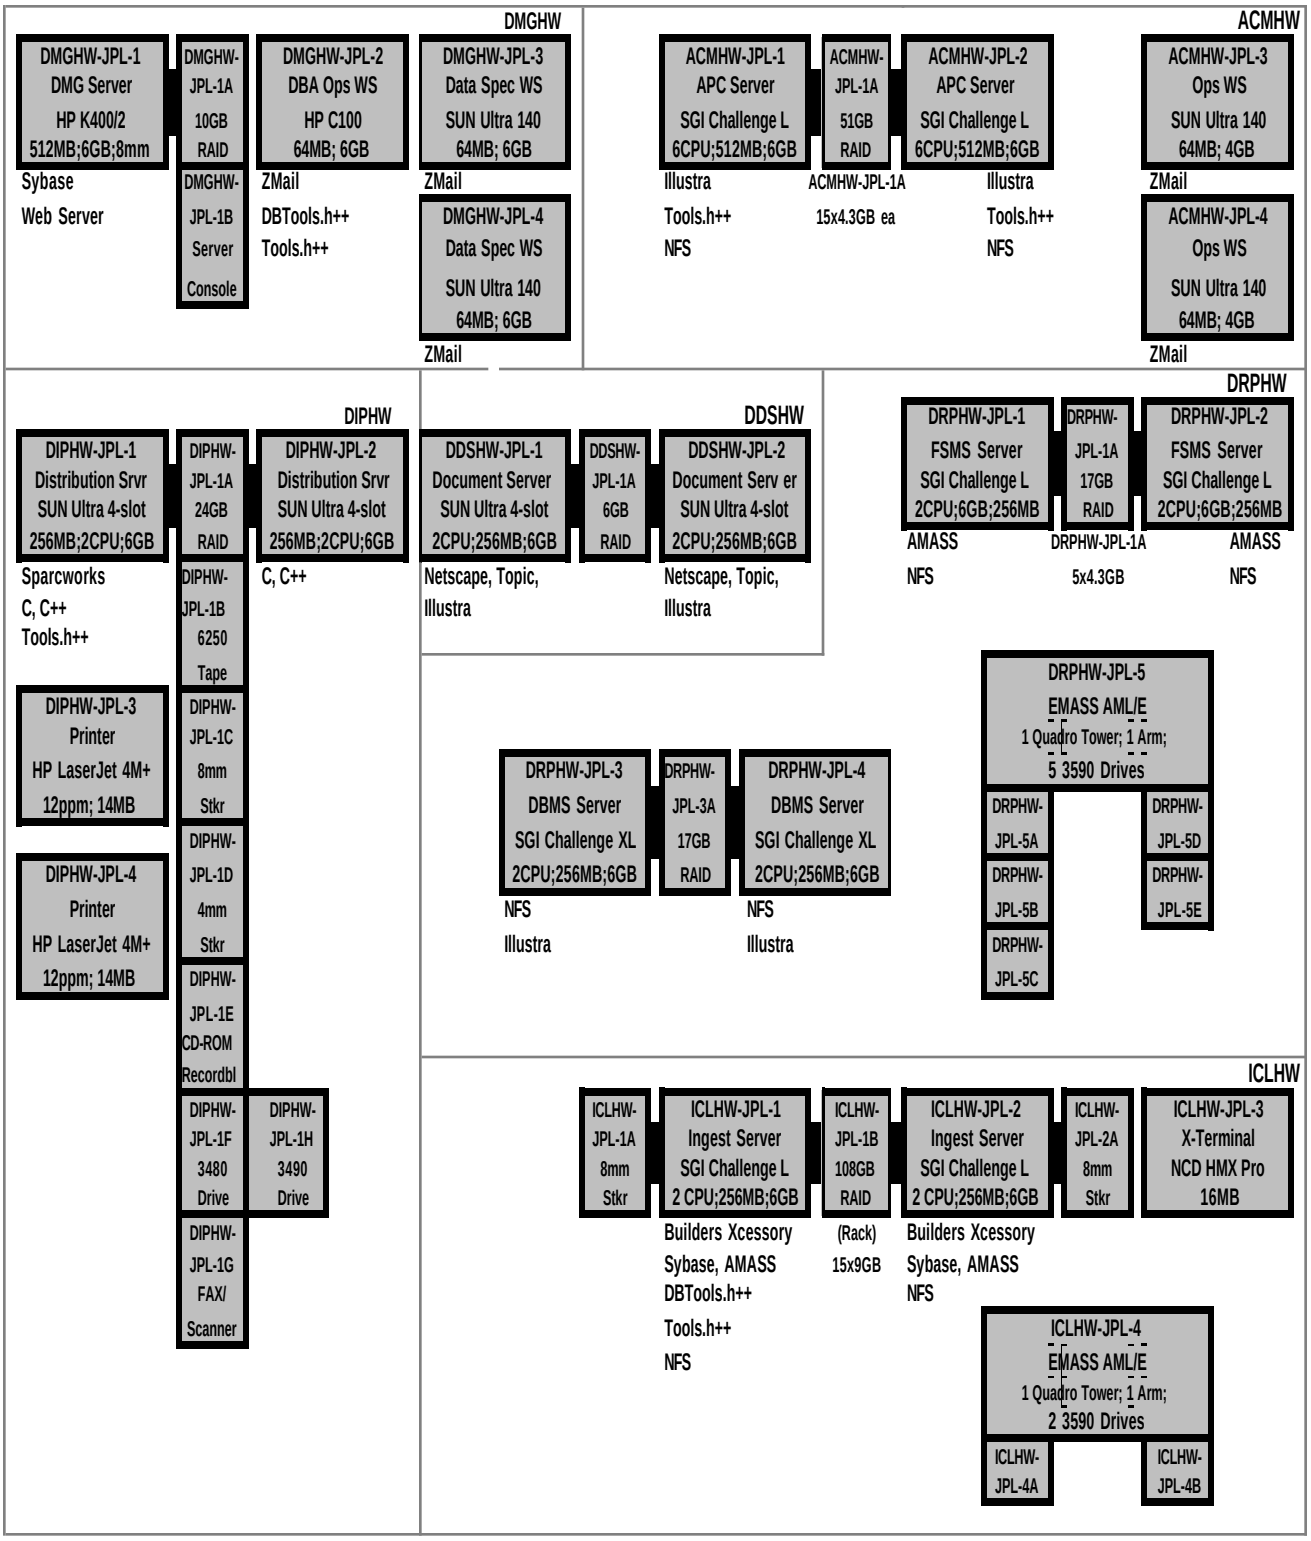

Figure 5-2. JPL ECS DAAC Hardware Configuration (2 of 2)

5.3 ORNL ECS DAAC

The ORNL ECS DAAC is illustrated in Figure 5-3, ORNL ECS DAAC Hardware Configuration. The ORNL ECS DAAC does not include production capabilities and will use storage capabilities already present at ORNL, and therefore includes hardware and COTS software for the following HWCIs only: MSS, CSS, ISS, DMGHW, ACMHW, DDSHW, and DRPHW. The Ingest software is hosted by the ACMHW CI at ORNL.

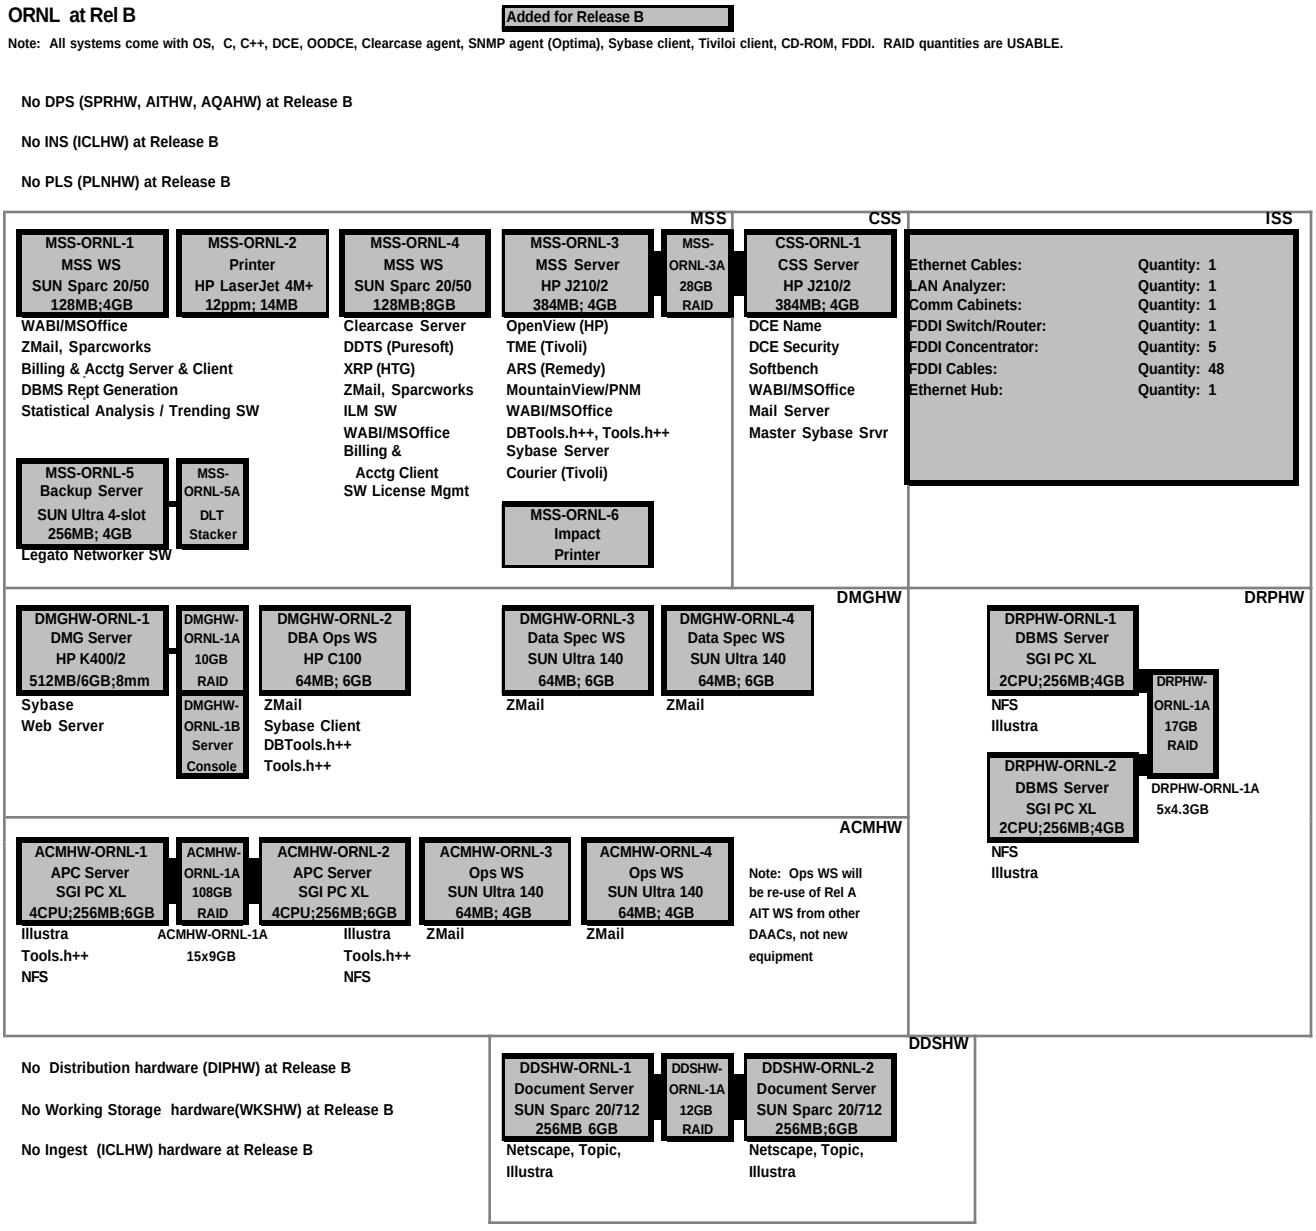

Figure 5-3. ORNL ECS DAAC Hardware Configuration

5.4 ASF ECS DAAC

The ASF ECS DAAC is illustrated in Figure 5-4, ASF ECS DAAC Hardware Configuration. Current plans for the ASF ECS DAAC are to include hardware and COTS software from the following subsystems only: MSS, CSS, ISS, DMGHW.

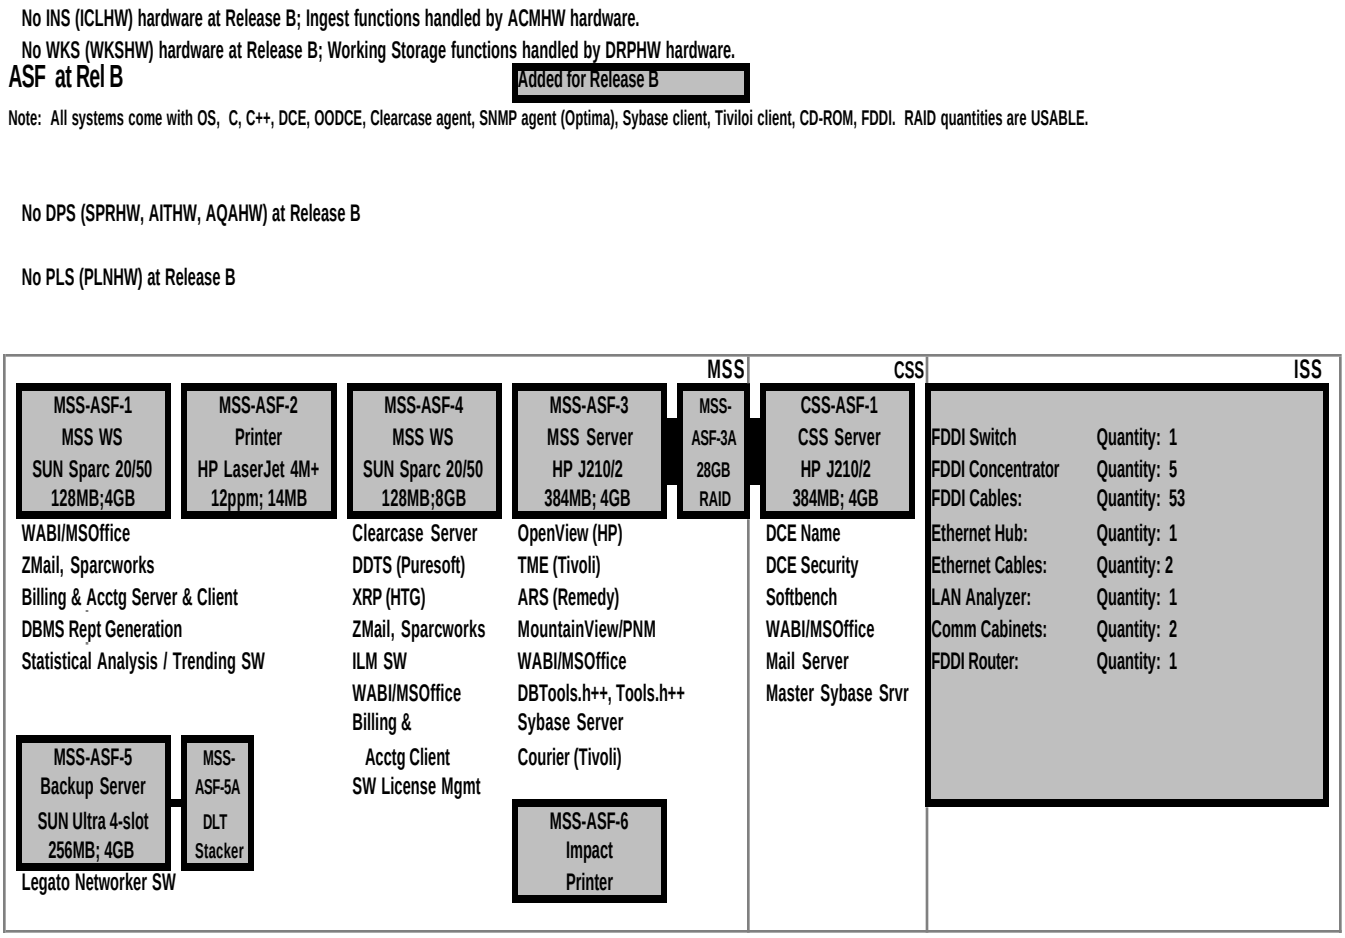

Figure 5-4. ASF ECS DAAC Hardware Configuration (1 of 2)

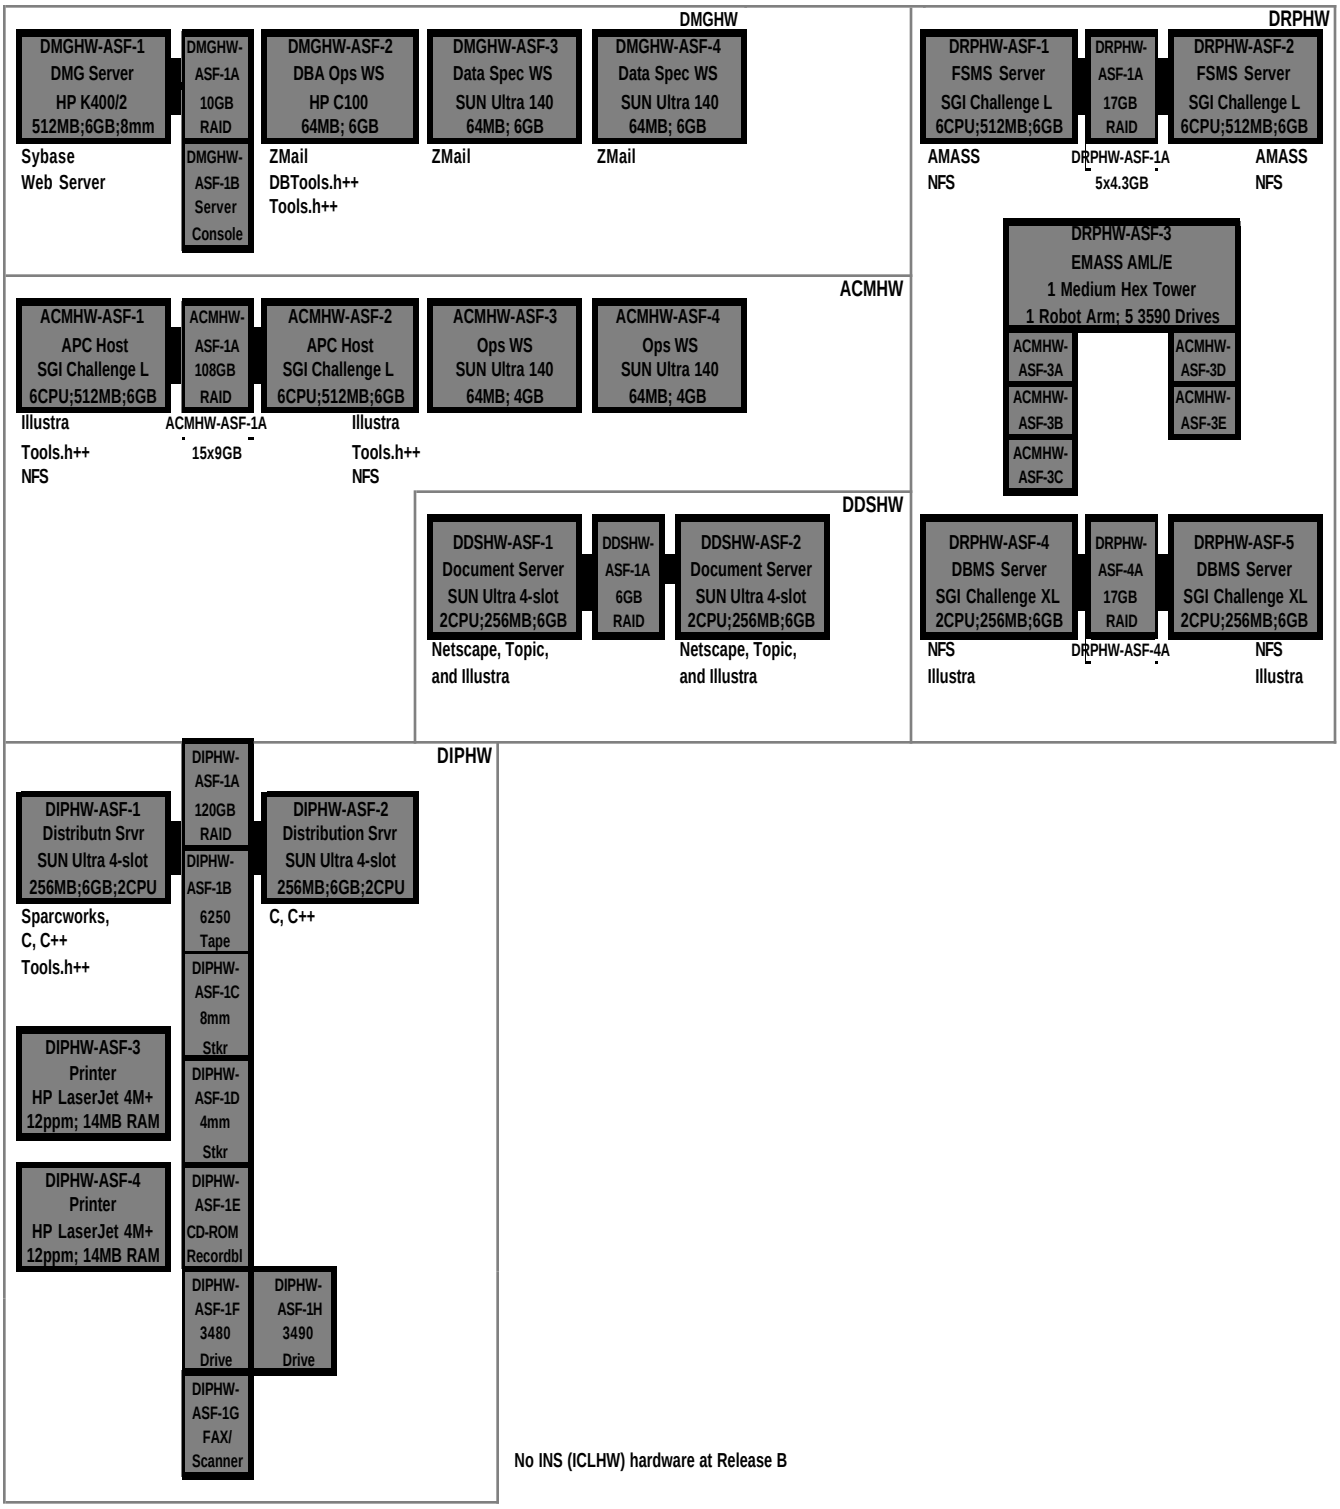

Figure 5-4. ASF ECS DAAC Hardware Configuration (2 of 2)

6. Program and Project Constraints and Assumptions

6.1 Constraints

1.≠ The system is not entirely partitionable into two independent systems for the purposes of parallel testing and operations. In particular the sites do not, in general, have fully redundant archive hardware.

2.≠ The TRMM mission will be operational, requiring significant operational usage of the system. If downtime is unavoidable then the critical concern is how quickly (and how surely) the operational system can be restarted. When downtime is recommended, the cost of the alternatives should be given.

There is a distinction between unstaffed hours, which vary from DAAC to DAAC, and actual system downtime, since the pull-side system remains available to users when the system is not staffed.

6.2 Assumptions

1. Release B hardware will be available at sites by CSR.
2. There is no requirement to run the Release A and B DBMS Servers in parallel.

This page intentionally left blank.

7. Identification of Transition Issues

7.1 Overall System Issues

These issues include configuring for test, piecemeal versus "big bang" transition, A/B site interoperability, DCE versioning, and site-developed software.

7.1.1 Test Configurations

How can Release B be tested with minimal impact to Release A operations? The sites where this issue arises are primarily SMC, GSFC and LaRC. EDC is not a production site in A, but is supporting interface testing and SSI&T.

The system can be shared in three ways:

1. Partitioning the hardware and network into distinct systems
2. Serializing use of the system so that either Release A or B, but not both together, use the system.
3. Using mode management to run the two releases in parallel.

The following options have been identified and the cost, feasibility, and mission impacts of each option are being assessed.

7.1.1.1- Option 1: Install B hardware, then run A in A environment OR run B in B environment

This serialized option is under investigation for use at the EDF. It has the highest downtime impact on A.

7.1.1.2- Option 2: Install B hardware and partition into separate A and B systems with no network connection to each other

Suboption 1 is to run separate A and B systems ECS-wide. B-only sites interoperate with each other but not with A sites.

Suboption 2 is a separate B system at the site (may need to stub interfaces to other sites)

How to provide two sets of Data Server robotics? Options are:

1. B system uses an archive simulator during testing
2. B system uses the archive robot that is added to Ingest in B in both Data Server and Ingest configurations

7.1.1.3- Option 3: Upgrade A to B Infrastructure, then run A in B environment OR run B in B environment.

This option would use the sustaining engineering process to install Release B versions of operating systems and DCE into Release A. It is similar to option 1 but with a faster switchover to and from test mode, since reconfiguration is much simpler.

7.1.1.4- Option 4: Upgrade A to B Infrastructure, then run A and B concurrently in the B environment.

Requires more (but not all) mode management functionality to be installed in Release A. Release A incorporates mode awareness in processes, so modal processes can run in test mode for B and in ops mode for A, but capability needs to be added to Release A to share multimodal processes (eg. archive, event logging).

This approach has the lowest downtime impact on A operations. However note that stress testing, failover testing, and some tests involving external interfaces cannot generally be run in parallel with operations. While these kinds of tests will take place at the EDF where possible, some further checkout at sites will be needed. Special measures could be introduced to support some tests, eg. external provider simulator to test external interfaces.

7.1.1.5 Test Configurations Trade-Off Analysis.

Options 1 and 3 require serial, rather than parallel, use of the system and therefore cause significant downtime for Release A. It is expected that, when system usage requirements for the operational missions are known, there will be unacceptable impacts. Option 1 is also impractical because it would take too long to change the framework when switching from A to B.

Option 2 partitions the system, but it is expected that there will not be enough capacity to make 2 representative systems, when equipment outages and other contenders for system use are known. However partitioning a small part of the system will be studied further, since it could provide a way to support, for example, SSI&T (SSI&T can operate without Data Server if data is pre∞ staged).

Option 4, parallel operations and test using mode management, is the preferred approach for Release B test activities that cannot be conducted using a partitioned system. Installation of the Release B framework and mode management into Release A, via a sustaining engineering release, is discussed further in section 8.1.

Note that more than one option will be used during transition. Stress testing Release B may need exclusive use of the system, while SSI&T may use a partitioned system, and routine tests may be accomplished via mode management.

7.1.2 Piecemeal Transition Versus "Big Bang"

"Piecemeal" means that Release A and B components will interoperate within the system. "Big Bang" means that either Release A executes or Release B executes, with no interoperation. There are two aspects to this issue:

1.≠ Piecemeal Testing - Can Release B components be tested against trusted Release A components?

2.≠ Piecemeal Operations - When testing is complete, should transition from A operations to B operations be conducted in a piecemeal fashion, component by component, or should all the Release A software be replaced with all the Release B software?

7.1.2.1 Piecemeal Transition

Piecemeal transition levies requirements on some Release B software to meet both A and B interfaces. Guidance has been given to Release B developers to maintain backward compatibility with Release A software, where possible. The guideline for APIs that have to support both A and B clients is "don't break the interface", or specifically:

- Methods can be added to objects but existing methods should not be changed or deleted
- Method parameters should not be changed or deleted
- ≠ Method parameters that are added for B should be at the end of the signature, and should be defaulted if not used by A.

These guidelines mean that, generally, service requests by A clients (in the formal, not end-user sense) will be honored by B servers; B client requests may or may not be understood by A servers.

An analysis of the subsystem interfaces (see subsection 6 for each subsystem in section 7.2) has indicated that some additional common infrastructure components would need to be installed to permit a piecemeal approach to work. Specifically, services such as those offered by MSS, CSS, and DSS would need to support the piecemeal mix of A and B clients. This would require a much larger release than that proposed in Section 8.1, because a piecemeal configuration requires most services to be compatible with both A and B clients, while mode management requires only a very limited number of multimode servers to be compatible with both A and B clients. Making the major assumption that such a common infrastructure could be put in place, some possible piecemeal configurations are as follows (A' denotes the Release A system as modified by the sustaining engineering release):

1.≠ Ingest-B, with rest of system at A' (would require stubbing of interfaces from Ingest to Data Server, but would allow Release B external interface testing to be conducted).

2.≠ Others TBD depending on how far Data Server-B is backward compatible with rest of system at A'.

7.1.2.2 "Big Bang" Transition

This strategy separates the Release A and Release B software, with no interoperability (other than the shared infrastructure). Test drivers would be used to test individual B subsystems. Once testing and transitional activities are complete, cutover to B would occur, with the option to switch back to Release A as desired. It may appear to be higher risk, when considered in isolation.

7.1.2.3 Piecemeal Versus "Big Bang" Trade-Off Analysis

Advantages of piecemeal interoperability are:

1. Transition is broken into a series of smaller, more easily manageable, lower risk pieces.
2. ≠ It may allow testing of Release B subsystems against existing, reliable Release A subsystems, and it may allow Release A and B sites to interoperate (however interoperability between sites is needed only for the very few Release A inter-site interfaces).
3. ≠ Sites will transition at different times, so some limited interoperability will be needed in any case.

Disadvantages are:

1. ≠ A common infrastructure to support the piecemeal mix of A and B components would need to be put in place. This could be unachievable since there may be significant changes in the implementation of common services from A to B, such as the addition of the Server request Framework to some subsystems. Even if it were technically feasible, it would require the addition of significant functionality to the proposed sustaining engineering release, which would create schedule and process problems.
2. ≠ Interfaces to Release A would have to be developed by Release B, and then thrown away once B is operational. Stubs would have to be developed for each unsatisfied interface.
3. ≠ Each piecemeal configuration used operationally would require full testing and certification, causing a large increase in overall test costs.

These disadvantages substantially outweigh the advantages. The "big bang" approach is the default, and avoids the costs of:

- maintaining backward compatibility in all common services
- testing each piecemeal configuration
- stubbing incomplete functions
- making several software releases.

Because the big bang approach is much cheaper and simpler, and because its risks can be managed, it is selected. However Release B interfaces will continue to be studied to look for low-cost ways to test B components against A components.

7.1.3 Interoperability Between B And A Sites

During transition, Release B sites may interoperate with Release A sites, levying requirements on release B to meet the cross-site interfaces presented by Release A. For example if SCFs can talk to multiple DAACs, they may see heterogeneous interfaces.

The current baseline performs IATO and IV&V activities at Release A sites before the B-only sites, but a better overall order may be be to install and test first at the (presumed less busy) Release B-only sites, and then at the Release A sites. Which site is first? Should EOC be late,

due to the sensitive nature of spacecraft control, or early since it does not go operational until post-transition? The operational Release A sites may be turbulent. SMC will be involved throughout, so must transition early. Some B-only sites (ASF and ORNL at least) have operational systems that may be impacted by the arrival of ECS. These issues need further study, followed by a revisit to the baselined site acceptance testing plan to rearrange the site order.

Based on the piecemeal versus big bang arguments above, a big bang transition of all sites to Release B operations is presumed. Operations involving mixed A and B sites would require not only additional test effort but also special procedures to avoid attempts to use Release B capabilities that are not yet present at the Release A sites. Furthermore, mixed A and B sites could force users to run two separate client sessions, if they wanted to access two sites.

7.1.4 DCE Client-Server Version Numbering

DCE-based clients and servers (essentially all custom ECS code) use major (j) and minor (i) version numbers in their Interface Definition Language (IDL) to specify their compatibility, using the form j.i. For example clients at 1.8 will be compatible with servers at major version 1 and minor version 8 or higher. A change in the major version number indicates that the server is no longer compatible and clients must change. A change in the minor version number indicates a small change that does not require clients to change, although if the change is the addition of a method or service the client may want to change to take advantage of it. Once added, IDL signatures cannot change without a major version number change, and IDL does not permit overloading of signatures.

During transition to B, some servers, for example the DCE CDS name server, must support both A and B clients. Other servers will exist in multiple instances; for example an instance of the Advertising Server normally exists at each ECS site, and during transition some instances may be Release A versions while others are Release B versions. Therefore where the Release B version is incompatible, or where compatible but interoperability is not required, Release B developers could use a different major version number to Release A, and where the Release B version is compatible and may interoperate with A, the major version number must be the same as A.

7.1.5 Site-Developed Software

During the operational lives of IR1 and Release A the sites will develop a legacy of scripts, configuration data (eg. OpenView maps), home pages for the web, and other site-specific software. They are likely to get quite concerned that the transition to B will break their software. This section will explain our strategy for addressing their concern, which is to:

1.≠ Provide guidance to the sites on how to write software that won't break during transition. The guidance will be developed in this technical paper. Currently known pitfalls are:

   a. Cell names, which will change during transition.

   b.≠ DCE-based client-server version numbers, which should be compatible where Release A and B will interoperate.

   c.≠ Mode awareness: this is needed when software connects to a resource (like a database or a server) to make sure that the connection is to the correct resource for the mode.

An example of site-developed software that needs to be mode-aware is a script that executes queries on a Sybase database; the script must connect to the database instance corresponding to the mode. The script could get the mode from a command∞ line argument, environment variable, or configuration file.

d. Others, TBD

2. Capture and back up site-developed software prior to transition.

3.≠ During transition, upgrade site-specific software as necessary. ECS's formal obligation is to upgrade only contract software deliverables; experience from the IR1 to A transition will indicate how much additional software is developed at the sites.

7.2 Subsystem-Specific Issues

The Release B subsystems are Client (CLS), Management Services (MSS), Planning (PLS), Processing (PRS), Interoperability (IOS), Data Server (DSS), Data Management (DMS), Communications And Infrastructure (CSS), Data Ingest (INS), and Internetworking (ISS).

7.2.1 Client Subsystem

The Client subsystem contains the Desktop and Workbench CSCIs.

7.2.1.1 Brief Summary Of Changes

Introduction of the Release B Client , but the Release A V0 client remains supported until phase out of the V0 system. Users are encouraged to install the Release B Client.

Interface from most Release B Clients to servers will be direct, not via the GTWAY CI (kerberized V1 may still use an enhanced GTWAY).

Replacement of the V0 based Earth Science Search Tool (ESST) with an ECS developed ESST. Requires support of the Earth Science Query Language (ESQL) throughout the system (biggest impact is Data Server; GTWAY also affected since ESQL is not same as in A).

User interface for Data Acquisition Requests (DARs) to schedule instrument data acquisition is added.

User interface for On-Demand Processing Requests (ODPRs) is added. Needs support from Data Server and PDPS; for certain products DPRs may be sent via ASTER and ASF gateways (also an indirect impact on GTWAY code because we provide it to the Japanese Aster team, and they use it for the incoming interface from ECS).

7.2.1.2 Assumptions And Dependencies

Desktop interfaces (e.g., the Ingest operator GUI) must not be upgraded until the corresponding capabilities are available in Ingest. The same is true for the Release B Client.

7.2.1.3 Data Changes

The Release A desktop has persistent data called Dependent Valids that is no longer needed by the Release B Client (because the Release B client is able to query the Data Dictionary).

Science users can save certain kinds of data, such as granules, from the desktop as files. DAAC users can save a wider set of data kinds. Further study of the desktop is needed to identify all kinds of data that could be stored, in order to identify any possible problems processing stored files following transition to B.

7.2.1.4 Hardware Changes

Not applicable.

7.2.1.5 COTS And Custom Code Changes

Common Desktop Environment (CDE) COTS, and Java, are being considered for the Client. Two prototype clients will be developed, with a selection made at the conclusion of EP7 in November 1996. One development will use established X/Motif technology while the other will use Netscape 2.0 and Java.

7.2.1.6 Interoperability With Release A Components

Release B Client generally needs the Release B Interoperability and Data Management susbsystems present (at a minimum) plus core support from MSS-B and IDG-B. However Release B application GUIs, such as the Ingest GUI, could be tested via a simple modification to the Release A desktop.

7.2.2 Management Services Subsystem (MSS)

The Management Services Subsystem contains the Management Agents, Management Logistics , and Management Software CSCIs, and the Management Hardware HWCI.

7.2.2.1 Brief Summary Of Changes

The changes from Release A to Release B are mostly upgrades to existing COTS packages, configuration changes to the existing COTS package, and the incorporation of new COTS packages. There are several assumptions and dependencies that, if not adhered to, will modify greatly the content of these sections.

7.2.2.2 Assumptions And Dependencies

Release A will upgrade HP OpenView to version 4.0.

Must upgrade from Sybase version 10.0 to version 11.0 of Sybase for the Billing and Accounting COTS package.

7.2.2.3 Data Changes

Currently there are very minimal planned schema changes from Release A to Release B.

7.2.2.4 Hardware Changes

The only change to the MSS server for Release B is the addition of 20GB of RAID to the memory shared with the CSS server.

7.2.2.5 COTS And Custom Code Changes

1. Agent and MIB changes to handle new applications.
2.≠ Interfaces from MSS to clients will change due to billing & accounting, and due to production events capability being added to end-to-end request tracking.
3.≠ MSS Performance Management: Add COTS Statistics package for trending. May require some additional custom code.
4.≠ MSS Fault Management: Configure TIVOLI (from Release A) to perform fault correlation.
5. MSS Security: Configure for multiple DCE Cells.
6.≠ MSS Accountability: Request tracking. Add small number of attributes to existing class objects.
7.≠ MSS Baseline Manager: If any, slight modifications for Release B. COTS package has already been selected for Release A (HTG-XRPII).
8.≠ MSS Software Change Manager: If any, slight modifications for Release B. COTS package has already been selected for Release A (ClearCase).
9.≠ MSS Change Request Manager: Slight modifications, if any, for Release B. COTS package has already been selected for Release A (Distributed Defect Tracking System).
10. Addition of Report Writer COTS.
11.≠ Security Management: Replacing HAL with custom GUI to native DCE interface when upgrading to DCE 1.1.
12. MSS Physical Configuration Management: Configure for additional DAACs.

7.2.2.6 Interoperability With Release A Components

In order to perform full Release B functionality, there are dependencies on Data Server and PDPS for new/modified interfaces. The interfaces can be dormant as long as the functionality is not executed.

Other areas of new/modified interfaces for MSS include Request Tracking and user profile. For User Profile, there will be additional attributes added to signatures. With Request Tracking, there are changes for Release B. It is TBD whether backward compatibility with Release A will be provided.

MSS interfaces with other ECS sites in both A and B for Trouble Ticketing, through email, and with the SMC for management reports.

7.2.3 Planning Subsystem

The Planning subsystem comprises the Production Planning CSCI and the Planning Hardware HWCI.

7.2.3.1 Brief Summary Of Changes

On-Demand Production Request support; this includes a new interface with DSS so DSS can send production requests to Planning. It may also require the Client subsystem to access the PDPS database or at least have access to Planning's PGE information. It is a new CSC, and a new process will exist to handle DSS requests.

Resource Planning enhancements; new types of resources, including knowledge of DSS and INS will be added. This shouldn't require any new system interfaces and is a relatively minor change.

Reprocessing Request support; this will require modifications to the Production Request Editor, and changes to the user interface. This change shouldn't require any new interfaces.

Production Rules; basically additional code in a number of CSCs and a new user interface to modify production rules for Production Requests. A number of new fields will be added to the PDPS database.

Production Strategies; Incorporates a new user interface to be able to do this. However there are no new subsystem interfaces.

7.2.3.2 Assumptions & Dependencies

Since PLS & DPS share the PDPS database as a communication mechanism, both will require the same database version, and therefore they will have to be using the same version (i.e., either Release A or Release B, viz. Sybase 10 or Sybase 11).

It is likely that PLS & DSS will need the same software version also. The interfaces for services such `as query, inspect, and acquire are changing slightly between Releases A & B. If these were to remain constant, PLS A could run with DSS B as long as DSS did not make On-Demand Production requests. Likewise, PLS B could run with DSS A as long as PLS did not request any metadata subscriptions.

The above constraint applies to the MSS and ADSRV processes also.

PLS and INS should be able to run in any combination since the only interface is through subscriptions.

7.2.3.3 Data Changes

See database comment above. A script will be run to convert the PDPS database from Release A data into Release B data format.

7.2.3.4 Hardware Upgrades That Affect Software

MRS is considering upgrading the PLS machine in a number of ways. If they include a port to a different box (i.e., from a SUN to an HP or SGI), then the custom code and COTS will need porting. Likewise, a dedicated PDPS database server may be incorporated in Release B (it is hosted on the same machine as the Planning Server in Release A), which would require modifications in the database code, but would ease transition since the new processor could be used for testing.

7.2.3.5 COTS Upgrades/Changes

A new AutoSys version may have some small effects on the PLS software.

7.2.3.6 Interoperability With Release A Components

See section 7.2.4.6.

7.2.4 Processing Subsystem

The Processing Subsystem comprises the Processing, Science Data Processing Toolkit, and Algorithm Integration And Test CSCIs, and the Science Processing, Algorithm Integration And Test, and Algorithm Quality Assurance HWCIs.

7.2.4.1 Brief Summary Of Changes

Predictive Staging. Some data may be staged before the specific PGEs are ready to execute to improve throughput of the system. This involves new fields in the PDPS database which will indicate if a PGE is to have its data predictively staged.

Reducing the number of jobs in the job box. For Release A, the Production Monitor will see (for each PGE) eight jobs on the AutoExpert display. In Release B that number is reduced to 3 or even 1. This will mean that when Release B is operating the Production Monitor will be able to actively watch fewer jobs through the AutoExpert displays.

Access to Science Software Archive Packages. In Release B, the Science Software tester will have the ability to modify, create, and delete Science Software Archive Packages (groupings of source code, test data, documentation, etc.)

Activation Rules. Release B will support more ways to activate PGEs than Release A. This requires changes to the PDPS database (added fields and new tables), and new GUIs in AITTL. Any existing PGEs should run as they did in Release A.

QA Metadata check. Quality Assurance will have the ability to check to compare the metadata of output products against expected values stored in the PDPS database. This function will require access to the Release B PDPS database, plus updates to the Release A PGEs so that new metadata values can be checked against pre-existing PGEs.

Regarding SSI&T: Before mode management is available, Release B algorithms can be tested at Release A DAACs provided all the test data is staged so there is no data server access.

Alternatively, data server could be accessed, provided ops staff are prepared to read and write test data files from the operational data server archive.

When mode management is available, it will be possible for a test mode to be established allowing data server access for an algorithm under test. The data server access would then go to the test partitions within the data server. The Resource Manager person must schedule the test as a ground event so that the operational system does not use the processor and disk.

7.2.4.2 Assumptions & Dependencies

For Science Software Archive Package updates, it is assumed that the Data Server will provide an interface in Release A.

It is assumed that both PLS & DPS will be the same Release (either both Release A or both Release B). You cannot run Release B DPS with Release A PLS.

7.2.4.3 Data Changes

The PDPS database will be changed extensively by the addition of new fields and new tables. It is impossible for Release B software to function with a Release A version of the database.

7.2.4.4 Hardware Upgrades That Affect Software

Any science processing machines that are added to the configuration need to be defined to both PLS & DPS so that they can be scheduled and utilized while processing science algorithms.

7.2.4.5 COTS Upgrades/Changes

COTS compiler upgrades for the SDP toolkit will cause Release A PGEs to be recompiled.

7.2.4.6 Interoperability With Release A Components

After the installation of the M&O.1 release (see section 8), PDPS-B can interoperate with Release A subsystems with no loss of Release A services provided core support from MSS-B and IDG-B is also present. In order to provide Release B services (eg. processing of new ESDTs), DSS-B must also be present.

PDPS has the following interfaces to other ECS sites in Release A: an X11 interface between the Science Computing Facility and the Production Request Editor user interface in the Planning subsystem. In Release B there will be interfaces to the Data Servers at various DAACs for Plans & Schedules and data granules.

7.2.5 Interoperability Subsystem

The Interoperability Subsystem contains the Advertising Service CSCI and the Advertising Service HWCI.

7.2.5.1 Brief Summary Of Changes

The Interoperability Subsystem (IOS) consists of one CSCI, the Advertising Service (ADSRV), which has two primary CSCs, the Advertising HTTP server and the Advertising Application Server. The HTTP server is COTS. It is assumed herein that it will be the same as the one used in Release A, but there will be an additional instance in Release B. The HTTP server COTS final selection for Release B is TBD. The ADSRV software runs on the DMS hardware CI, so at Release B the ADSRV HTTP server will be installed, configured, and started on the DMS hardware CI.

There are two primary functional capabilities added by Release B. These are:

1.≠ Support for subscriptions on the insertion, deletion, and update of advertisements. It is assumed that the ADSRV will notify the subscription server resident on the DSS HW of these events and the subscription server will process the subscription.

2.≠ Integration with the Data Dictionary Service (DDICT) of the DMS to provide definitions of the terms in the advertising service. This integration will occur in the ADSRV application run from the HTTP server and it will call up the DDICT web interface serviced by the same HTTP server.

7.2.5.2 Assumptions & Dependencies

HTTP server COTS selection remains the same as in Release A.

The ADSRV will share the Sybase SQL server and Replication Server installation with the DMS. Therefore, it will follow the same release versions as the DMS, Sybase 10 for Release A and Sybase 11 for Release B. The data will be copied from the Sybase 10 database to the Sybase 11 installation using the same mechanism as the GTWAY data.

7.2.5.3 Data Changes

The Release B ADSRV database will primarily be of the same structure as the Release A database. Any minor modifications will be accomplished through SQL scripts. The SQL scripts will initially create the Release B database. Next the data will be copied from the Release A database using a second set of SQL scripts. The Release B SQL server will be a trusted server of the Release A SQL server, so data can be exchanged through this mechanism.

7.2.5.4 Hardware Upgrades That Affect Software

None of the hardware changes occuring with the Data Management hardware will affect the software installation or configuration of the ADSRV.

7.2.5.5 COTS Upgrades/Changes

Release A will be using Sybase SQL Server 10.0.1. Release B will use version 11. Since the Release A software will not be upgraded to access the Sybase 11 database, the Sybase versions will be installed in parallel. The Release A data required for Release B will be copied via SQL scripts and by making the Release B SQL Server be a trusted server of the Release A SQL

server. The Replication Server will also be upgraded and installed in parallel with the Release A instance. The Release B ADSRV database will be loaded and replicated to each DAAC using the Replication Server. Once the Release B system becomes operational, the Release A Sybase SQL server and Replication Server can be shut down and the data removed.

An HTTP server will also be installed to service Web clients of the ADSRV, Client Subsystem, and the DDICT. It is assumed that this is the same HTTP server used in Release A, but a new instance will be installed on the DMS hardware instead of sharing one with the Document Data Server. The Release B URL to the ADSRV will have to be made known to the community, since it will likely be different than the Release A URL. A pointer from the old URL to the new one will have to be maintained for those users who do not get the message.

7.2.5.6 Interoperability With Release A Components

IOS will not interoperate with release A components (except the Release A Client via the V0 Gateway, of course). IOS requires the Server Request Framework and the Rel B version of MPF. Assuming that there will be some change to the DSS client libraries from A to B, the IOS components will not talk to the DSS-A components.

7.2.6 Data Server Subsystem

The Data Server Subsystem contains the Science Data Server, Document Data Server, Storage Management, and Data Distribution CSCIs, and the Access And Control Management, Working Storage, Data Repository, and Distribution And Ingest Peripheral Management HWCIs.

7.2.6.1 Brief Summary Of Changes

Science Data Server: Data Type Services: On-Demand Processing and Compound; Session & Request Management: Subsetting, Suspend and Resume Sessions, and Accounting; Data Type: Metadata Problem Reports, Versioning, and Special Data Products; DARS: All new DAR Requirements; Subscriptions: Time Intervals.

Storage Management: Reporting of Disk Utilization Information, Estimation of DSS Resource Utilization, Deleting files from archive, Statistical Monitoring of checksum error rate, Archive data backup and restore. Separate working storage server: separate from archive servers and multiple servers can be provided if necessary. Implement hierarchical working storage using RAID and fast robotic library. Second tier of temporary storage for managing network ingest and distribution. Separate networks to minimize I/O contention: working storage and archive and working storage and data processing. Introduction of bulk data service between working storage and data processing. Security “firewall” prevents user access to data repository.

Data Distribution: Cost estimating, Cost accounting, Alert SMC about electronic transmission problems, Distribute on 3480/3490, Distribute on facsimile, Automatic generation of media labels, automatic generation of shipping labels, and compression.

Document Data Server: Possible migration of document metadata and keywords to DDSRV's Illustra database, which may in turn affect the sizing for the DDSRV DBMS machine. In

addition, management of DDSRV will be improved and more document formats, including WordPerfect and Interleaf, will be supported.

7.2.6.2 Assumptions And Dependencies

Assume no changes to the Network Infrastructure and that protocols for B are the same as A except in the database area.

Data Server will receive ESQL as well as Parameter List inputs from the Earth Science Search Tool in Release B.

The COTS DBMS for the Document Data Server is assumed to remain Illustra, but the actual selection is TBD.

7.2.6.3 Data Changes

Release B adds about 180 new data products. These products are listed in appendix F of DID 304. "Data" covers product (files) data, metadata, and system infrastructure data. When science data is migrated we must maintain synchronization between the archives and the metadata, as well as system infrastructure data that will allow for continued location and retrieval.

The approach taken must minimize the effect of making URs obsolete. That is, if a client of data server has been given a UR, that UR should not be obsoleted by the migration of the data referenced by the UR from a Release A schema to a Release B schema.

7.2.6.4 Hardware Upgrades That Affect Software

Migrating from 3590 to D3 tape drives.

Adding facsimile capability.

Archive upgrades at some sites.

7.2.6.5 COTS Upgrades/Changes

IRIX 6.2, DCE 1.1, and Illustra, Open Client (multi-threaded open client solution for Sybase), Xalib for Sybase.

Version 11 of Sybase needed for Release B versus version 10 for Release A.

7.2.6.5.1 Sybase to Illustra Migration

During the transition from Release A to B, all of the Science Data Server metadata stored in the Sybase DBMS will be migrated to the Illustra DBMS. The transition plan makes the following assumptions about the migration:

a..≠There is (currently) no requirement to run the Release A and B Science Data Server DBMS Servers in parallel.

b.≠ Based on the Release A DAAC-specific design documents, each DAAC will have less than one gigabyte of metadata to be loaded from Sybase into Illustra. This data includes

Release A V0 migration data and all new collection, granule, production history, document metadata, and ESDTs.

c.≠ The migration of Release B V0 data into Illustra during the Release B timeframe is considered a separate transition.

d.≠ There will be few data type changes or field length changes between the Sybase and Illustra schemas with little "data cleanup" required as in migrations from legacy systems to new systems.

e.≠ Document metadata and keywords will move from Sybase to Illustra, but will move into the Illustra DBMS on the Document Data Server, rather than the Science Data Server, in order to improve performance and ease future software development.

Because the Release A and B DBMS Servers do not need to run in parallel, the migration from A to B will be handled as an "initial database load," where at the start of Release B the Sybase DBMS is turned off and the Illustra DBMS is turned on. The initial database load will be accomplished in two basic stages. During the first stage, a series of flat files will be constructed from the Sybase database that contain all of the data currently stored in the DBMS as of a certain timestamp. The data in those flat files is then loaded into Illustra. This stage will take the majority of the load time because it will load most of the metadata. However, between the time that the flat files are built and the time that Illustra is turned on as the Release B database, subsequent updates to the Sybase database must be captured as the second stage load. Those updates must be captured during a small downtime into the same set of flat files generated during the first stage. These updates are then loaded into Illustra using a slightly different load method. The timeframes for the stages must be chosen in order to minimize the volume of data that needs to be loaded in the second stage. It is expected that the Sybase/Illustra benchmark prototype will provide sufficient information about load times to assist in establishing target dates for stage one and stage two.

The first step in effecting the generation of the flat file initial load data is to performing a mapping between the Sybase and the Illustra schemas. Although Illustra supports use of a completely relational schema, ECS is adapting the schema to take advantage of some of the object-relational Illustra features in order to enhance performance and to position the schema for future evolution during Release C. There will be a complete mapping between Sybase and Illustra; no data will be "eliminated" during the transition except the object id that is currently explicitly stored in Sybase as part of the ECS design. The object ID will be generated automatically by Illustra when the data is stored.

The second step is to develop the software that reads the Sybase database and generates a set of flat files, where each flat file contains the data that will be used to load a particular Illustra table. The software will be able to read the Sybase database containing the metadata while ECS is up and running. When the software begins, it reads (as a parameter) a timestamp that serves as the cutoff point for selecting updates so that the data written to the flat files is consistent. For example, if the flat file generation started at 5:00am on July 30, then the selected timestamp cutoff might be 4:55am on July 30. Any updates logged after the cutoff time would be captured in the second stage. This approach assumes that ECS software is treating referential integrity

correctly by capturing all necessary data relations between the start and commit of a transaction. Otherwise, the system has to be brought down in order to capture a consistent database.

As long as the object IDs that were generated and stored in ECS are maintained in the flat files, the flat files could be generated sequentially. That way, if the Sybase DBMS came down, only the file that was in progress would have to be regenerated (in the worst case).

Once the flat files are generated, there are three options that can be considered for the actual load of the data into Illustra. The Illustra bulk load process can be used, but there are currently known memory leaks in this that make this a poor choice. Normal Illustra database inserts can be used, but this is the slowest way to load the database for an initial load. The optimal choice for loading the data into Illustra is to use the direct access method. Direct access permits loading of the database through the Illustra engine (as in normal inserts), but bypasses the query planner and the checking pieces of the engine. Processing of the insert begins at the storage manager, which speeds up the load time. This approach places the burden of data correctness on the data in Sybase and the quality of data in our flat files.

One of the benefits of using direct access over the bulk load is that independent tables, i.e., they have no foreign key relationships, can be processed in parallel. In this way, the data can be loaded faster. If the Illustra DBMS comes down during the initial load, the load can be continued at the first unprocessed transaction.

During the second stage load, updates may have occurred to objects that already exist in the Illustra database. For these cases, the appropriate rows in the Illustra database will be deleted and the (new) updates will be inserted. New objects from Sybase behave exactly like they do in stage one.

Once the Illustra database is completely loaded after stage two, utilities must be run against both the Sybase and Illustra databases to ensure that all objects have been loaded into the databases. Row counts should be established and compared to ensure that all relationships have been maintained.

Although there is no requirement to keep the new Illustra metadata repository synchronized with the old Sybase repository, it is possible to establish a COTS architecture that would keep the two synchronized in order to revert back, if necessary. Illustra has an "Illustra to Sybase Gateway" that permits updates from Illustra to be sent to Sybase. Sybase also offers the OpenServer library that would permit Sybase updates to be copied to the Illustra DBMS. Neither product is part of the baseline, and would therefore be a COTS impact. A cheaper method of parallel maintenance would be to create and periodically run a script to extract updates from Illustra and copy them to Sybase.

7.2.6.6 Interoperability With Release A Components

TBD.

7.2.7 Data Management Subsystem

The Data Management Subsystem contains the Local Information Manager, Distributed Information Manager, Data Dictionary, and V0 Interoperability CSCIs, and the Data Management HWCI.

7.2.7.1 Brief Summary Of Changes

Release B includes 3 new CSCIs and the upgrade of one CSCI. The capabilities of these are as follows:

1.≠ The Data Dictionary Service (DDICT) is a CSCI that provides definitions of data collections, their attributes, domain values of attributes, and mapping of attributes and values to collection specific terms. This CSCI is used by the Data Server as a repository for a conceptual schema of the holdings of the Data Server. The Client Subsystem submits search requests to DDICT on behalf of user requests and for configuring the user interfaces. The DDICT database is used by the LIMGR, DIMGR, and GTWAY CSCIs to describe the schemas accessible through these components. The database is stored in Sybase and changes are replicated to each DAAC using the Sybase Replication Server.

2.≠ The Local Information Manager (LIMGR) is a CSCI that provides access to the services of a site including service requests to the GTWAY, SDSRV, and DDSRV. The LIMGR reads the DDICT database directly through a class library (RogueWave DBtools) that interacts directly with Sybase.

3.≠ The Distributed Information Manager (DIMGR) is a CSCI that provides access to the services across sites including service requests to the LIMGR, GTWAY, and SDSRV. The DIMGR reads the DDICT database through the DBtools library just as the LIMGR does.

4.≠ The V0 Gateway (GTWAY) is a CSCI that provides interoperability between V1 and V0. In Release A, the Gateway is from the V0 client to the V1 Data Servers only. In Release B, the interoperability becomes both ways; from V0 to V1 and from V1 to V0. In addition, the server will be upgraded to support all the mode management and event capturing requirements of Release B.

Each of the CSCIs in the DMS provide mode management capabilities as well as event logging consistent with Release B requirements. The LIMGR, DIMGR, and GTWAY CSCIs all access services in the DSS, using the Release B DSS interface.

Each Release A DAAC has two DMS HW servers which are HP K200 series server class machines. At Release B, the second K200 will be shifted to another DAAC and a CPU and RAM will be added to the remaining server. In other words, at Release B, each DAAC will have one K200 server with 2 CPUs and additional RAM. This provides comparable performance while decreasing the cost of two server machines at each DAAC.

7.2.7.2 Assumptions & Dependencies

The current baseline is that Release A will be using Sybase SQL Server 10.0.1 and Release B will be using version 11. As a result, the data from Release A that will be reused in Release B (only a small amount of data supporting the GTWAY), will have to be copied into the Release B DDICT database. (See the next section for more details on the data transition and the COTS Upgrades section on the Sybase transition.)

7.2.7.3 Data Changes

The GTWAY in Release A has a very small database to support the V0 to V1 interoperability. The schema of the DDICT database is much larger than the GTWAY schema in A. For Release B, the information used by the GTWAY will be stored in the DDICT database in a different table representation. SQL scripts will be used to copy the data from the Release A database to the Release B DDICT database.

7.2.7.4 Hardware Upgrades That Affect Software

The hardware upgrades at the sites between Release A and B should have no effect on the software.

7.2.7.5 COTS Upgrades/Changes

Release A will be using Sybase SQL Server 10.0.1. Release B will use version 11. Since the Release A software will not be upgraded to access the Sybase 11 database, the Sybase versions will be installed in parallel. The Release A data required for Release B will be copied via SQL scripts and by making the Release B SQL Server be a trusted server of the Release A SQL server. The Replication Server will also be upgraded and installed in parallel. The Release B DDICT database will be loaded and replicated to each DAAC using the Replication Server. Once the initial DDICT database is up and replicated, the configurations of the LIMGRs, GTWAYs, and DIMGRs can be performed at each DAAC. The changes to the DDICT database will be replicated since the replication environment will already be in place. Once the Release B system becomes operational, the Release A Sybase SQL server and Replication Server can be shut down and the data removed.

An HTTP server will also be installed to service Web clients of the Client Subsystem and to interact with the DDICT. It is assumed that this is the same HTTP server product used in Release A, but a new instance will be installed on the DMS hardware instead of sharing one with the Document Data Server. The HTTP server is actually a CSC of the Interoperability Subsystem, but it is installed and shared with the DMS.

7.2.7.6 Interoperability With Release A Components

None of the DMS CIs will interoperate with release A components. The exception might be that the Release B GTWAY will talk to the V0 IMS servers at Release A, but the Rel B GTWAY requires the Rel B client or some stubs. DMS requires the Server Request Framework and the Rel B version of MPF. Assuming that there will be some change to the DSS client libraries from A to B, the DMS components will not talk to the DSS-A components.

7.2.8 Communications And Infrastructure Subsystem (CSS)

The Communications And Infrastructure Subsystem contains the Distributed Computing CSCI and the Distributed Communications HWCI.

7.2.8.1 Brief Summary Of Changes

New: ASTER communications gateway (billing and accounting, security, DAR and DPR support, DCE cell-to-cell interface), and ASF communications gateway.

Distributed object framework upgrades as detailed in table 8-1.

Support for the Distributed File System (DFS).

Summary of relevant enhancements from DCE 1.0.3 to 1.1:

- •≠ Single Administrative DCE Control Program (dcecp) - provides a consolidated user interface across DCE, replacing the HAL tool from Release A.
- •≠ Serviceability Improvements - enhanced diagnostic messages. A new document, the *DCE Problem Determination Guide*, provides explanation and action to be taken for every DCE error code.
- • Cell Aliasing - enables cell names to be changed.
- •≠ Hierarchical cells - enables cells to be organized to match the hierarchical structure of an organization.
- •≠ Security Delegation - allows intermediary servers to operate on behalf of the initiating client across chained RPC operations.
- •≠ Extended Login Capabilities - includes password management, pre-authentication, and access only from trusted machines.
- • Performance improvements - including optimized RPC calls.

7.2.8.2 Assumptions & Dependencies

See table 8-1 for COTS version dependencies.

7.2.8.3 Data Changes

The name server and security server databases will be distributed among the sites when the multicell topology is established. For further discussion see section 9.1.

7.2.8.4 Hardware Upgrades That Affect Software

The only change to the CSS server for B is the addition of 20GB of RAID to the memory shared with the MSS server.

7.2.8.5 COTS Upgrades/Changes

See table 8-1 and section 9.1 for details of DCE and operating system upgrades.

7.2.8.6 Interoperability With Release A Components

It is expected that Release A (DCE 1.0.3) clients will interact with Release B (DCE 1.1) servers. The 1.0.3 to 1.1 prototype seeks to verify this.

The Release A Name Service will incorporate a configuration parameter to control whether Release A (single cell) or Release B (multi-cell) names are used.

7.2.9 Data Ingest Subsystem

The Ingest Subsystem contains the Ingest Services CSCI and the Ingest Client HWCI.

7.2.9.1 Brief Summary Of Changes

Release B includes the upgrade of the single software CSCI (INGST). The capabilities are as follows:

1.≠ The INGST CSCI provides the capability to control (hold, resume, and change priority of) ongoing ingest requests. The CSCI includes the operator GUI capability to specify requests to be controlled and the action to be taken.

2.≠ The INGST CSCI is restructured to fully distribute ingest request processing. In particular, a new Granule Preprocessing Manager program is deployed on each ECS processor on which ingest data preprocessing is performed.

3.≠ The INGST CSCI adds table information describing Release B data types. The table information allows table-driven metadata extraction, data conversion, and data reformatting.

4.≠ The INGST CSCI performs TBD data preprocessing (e.g., AMSR L2 preprocessing supplied by SeaWinds, CEOS conversions to EOS-HDF for ASF data) related to Release B data types.

5.≠ The INGST CSCI provides a GUI media ingest capability that provides access to new media types supplied at Release B. Specifically, TBD media is provided by ASTER at the EDC DAAC. Note: the media peripheral hardware and associated access software is provided by the DIPHW CI of the Data Server subsystem.

6.≠ The MSS startup script is updated to automatically start polling ingest daemons for new ECS polling with delivery record interfaces (e.g., EDOS at LaRC and GSFC).

Release B includes the upgrade of the single hardware CI (ICLHW). The capabilities are as follows:

1.≠ At Release A LaRC has a primary/backup pair of SGI Challenge L processors to support TRMM Level 0 data processing. No processor upgrade is expected at LaRC. No additional magnetic disk space (e.g., RAID) is added at LaRC. Instead, the current RAID storage is used for combined SDPF/EDOS working storage. CERES Level 0 data accumulated on RAID after TRMM launch is migrated to a new archive tape library unit.

4.≠ Hardware equivalent to that at the LaRC DAAC is installed at the GSFC DAAC to support EDOS Level 0 ingest. That hardware includes an SGI Challenge L pair, RAID similar in capacity to the LaRC DAAC, and an archive tape library.

5.≠ Hardware equivalent to that at the LaRC DAAC is installed at the EDC DAAC to support Landsat 7 L0R ingest. That hardware includes an SGI Challenge L pair, RAID similar in capacity to the LaRC DAAC, and an archive tape library.

6.≠ Hardware equivalent to that at the LaRC DAAC is installed at the JPL DAAC to support SeaWinds Level 0, ACRIM, and SAGE-III Level 0 ingest. That hardware includes an SGI Challenge L pair, RAID with TBD capacity, and an archive tape library.

7.2.9.2 Assumptions & Dependencies

The current baseline is that Release A will be using Sybase SQL Server 10.0.1 and Release B will be using version 11. As a result, the data type/metadata information from Release A that will be reused in Release B is copied into the Release B database. (See the next section for more details on the data transition and the COTS Upgrades section on the Sybase transition.) This is a software installation activity.

The MSS User Profile is updated to contain Release B external data provider information. This can be set up prior to the installation of general Release B software, so long as the structure of the User Profile does not change for Release B.

At initial system startup, the advertising service is accessed to determine available Release B Science Data Servers with insert services for Release B products. The advertising service is accessed by means of a custom operator tool. The advertising information for the Release B Science Data Servers must be updated prior to invocation of the operator tool.

The CSS Ingest Gateway (TCP/IP-to-OODCE gateway) must be configured to invoke the Release B version of the Ingest Server.

7.2.9.3 Data Changes

As mentioned above, new data type/metadata information is added for Release B data types. No data structure changes are required. SQL scripts will be used to copy the data from the Release A database to the Release B database.

7.2.9.4 Hardware Upgrades That Affect Software

The hardware upgrades at the sites between Release A and B have no effect on the software. Software to access and control new archive tape libraries is provided by the Data Server subsystem. Software to access and control new media peripheral hardware is provided by the Data Server subsystem.

7.2.9.5 COTS Upgrades/Changes

Release A uses Sybase SQL Server 10.0.1. Release B uses version 11. Since the Release A software is not upgraded to access the Sybase 11 database, the Sybase installations are installed

in parallel. The Release A data required for Release B is copied via SQL scripts and by making the Release B SQL Server a trusted server of the Release A SQL server. Once the Release B system becomes operational, the Release A Sybase SQL server can be shut down and the data removed (after appropriate backups have been taken).

Ingest HTML forms continue to access the Document Data Server HTTP server at Release B. Assumption--no impact to existing HTML forms nor associated script files.

7.2.9.6 Interoperability With Release A Components

Ingest, given core support from MSS-B and IDG-B, can operate with release A subsystems provided the interfaces to Data Server are stubbed. This would allow early interface testing for interfaces that are new with Release B. With the addition of Release B Data Server, Ingest could be fully tested.

7.2.9.7 Test Aspects

There are at least two options for providing test data inputs to Ingest. The first is to duplicate incoming inputs and direct one copy to operations (Release A) and the other copy to test (Release B). The second option is to record incoming data and replay it at a future time. These options need further study.

7.2.10 Internetworking Subsystem

The Release B Internetworking Subsystem (ISS) contains one CSCI, the Internetworking CSCI. This CSCI provides internetworking services based on protocols and standards corresponding to the lower four layers of the Open Systems Interconnection (OSI) reference model: the transport layer TCP and UDP protocols; the network layer IP protocol; and the physical/data link Ethernet, FDDI, and HiPPI protocols.

The ISS also contains one HWCI that provides the networking hardware for the intra-DAAC, DAAC to V0, DAAC to EBnet, SMC, and EOC connectivity, including: FDDI switches, concentrators and cabling, Ethernet/FDDI routers, hubs and cabling, HiPPI switches and cabling, and network test equipment.

7.2.10.1 Brief Summary of ISS Networking Changes from Release A to Release B

Release B's network design, for the most part, is a common design with similar topologies among the three sites - GSFC, LaRC, and EDC - transitioning from Release A to B. This design is shown in Figure 7.2.2-1 and it resulted from the following design changes to the Release A version:

- Add bandwidth to the DAAC and inter-DAAC networks to accommodate increased data volume.
- Implement a separate "User FDDI network" and "Production FDDI network."

- Replace the ECS network's direct connection to the campus's exchange LAN with a dedicated router connection. A new ECS router will be added for this purpose.
- Add an isolated, high-performance, point-to-point communications between the Data Server and Science Processor Subsystem hosts
- Modify various hosts network connectivity such that they become connected to multiple FDDI rings and HiPPI fabric.
- Add concentrator units to accommodate the additional single attached hosts and multi∞ homed hosts.

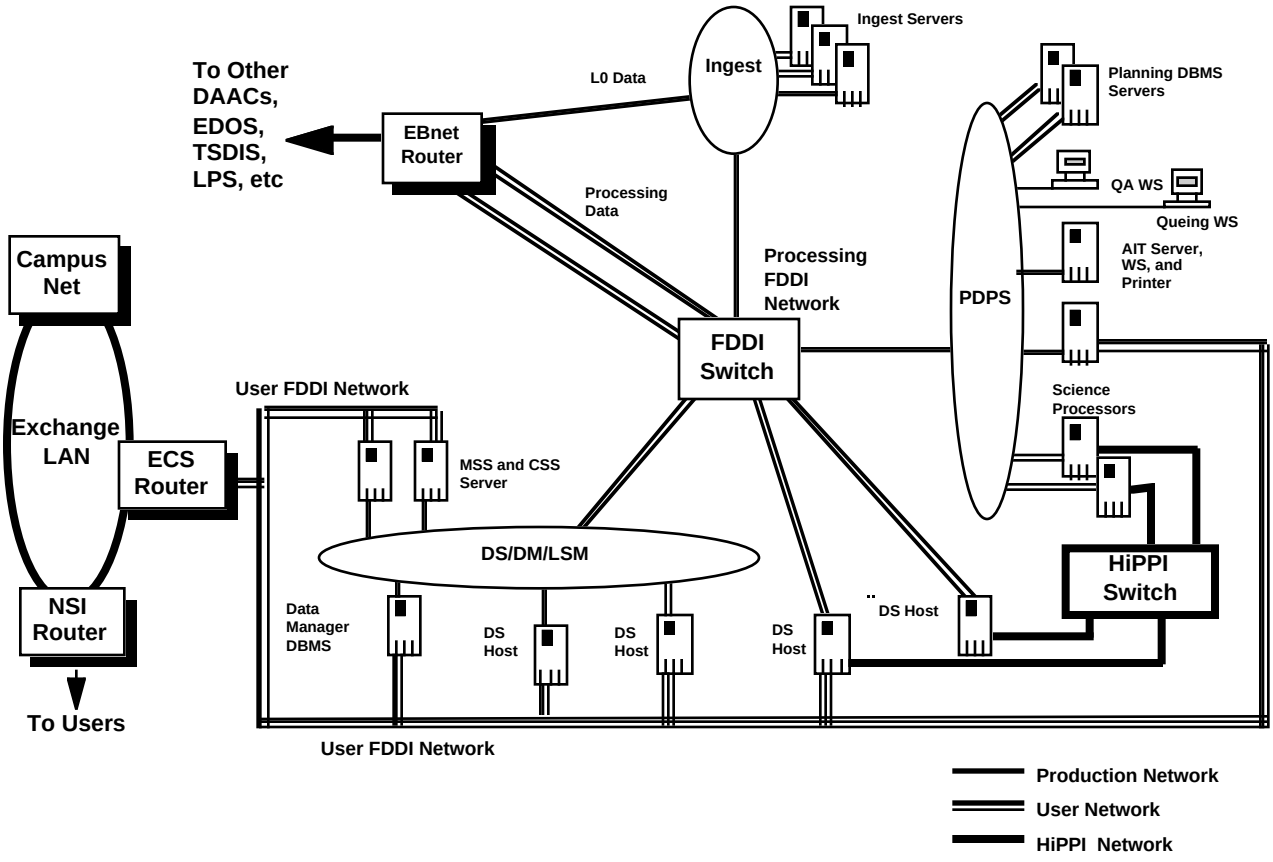

Figure 7.2.2-1 Generic Release B Network Architecture

7.2.10.2 DAAC-Specific Changes from Release A to Release B

Changes between Release A and Release B networks for GSFC, LaRC, and EDC DAACs are summarized in the following paragraphs.

7.2.10.2.1 GSFC ECS DAAC Network

The topology of GSFC's ECS DAAC Release B network design is quite similar to the Release A as-built network. Therefore, the topology changes necessary to transition from Release A to Release B are not viewed as major. Figure 7.2.3.1-1 shows the installed Release A network. Figure 7.2.3.1-2 shows the new Release B design. Summarizing the changes to be made:

- A new User FDDI network (FDDI ring) is added and specific Data Server Subsystem hosts will have an additional FDDI connection to this ring. Other Data Server hosts are attached directly to the FDDI Switch.
- The Release A LSM FDDI ring is merged into the Production ring. LSM subsystem (MSS, CSS) hosts are attached to the DM/DS/LSM FDDI ring and also to the User FDDI network.
- An ECS router is added to facilitate a routed connection to the GSFC campus's IsoLAN.
- FDDI Switch connections to the SMC and BBN subsystems are moved to the new ECS router connected to the IsoLAN.
- A high-performance, point-to-point, switched network using HiPPI technology is added to connect selected Data Server and Science Processor hosts for the exchange of very high data volumes. This HiPPI network includes a HiPPI switch.
- The EBnet WAN bandwidth is increased and a parallel FDDI connection added between the EBnet router and ECS FDDI Switch to add bandwidth needed for inter-DAAC data traffic.
- The EBnet router is connected directly to the GSFC IsoLAN and to the ECS router.
- The GSFC DAO network is linked to the ECS network's FDDI Switch (TBD).
- Where necessary, additional FDDI concentrator units and/or ports are added to accommodate the additional hosts required for Release B.

The Release A IP network and host addressing assignments were selected with the Release B network design in mind; therefore, it is anticipated that no network numbering changes will be required and few, if any, Release A hosts addresses will be changed.

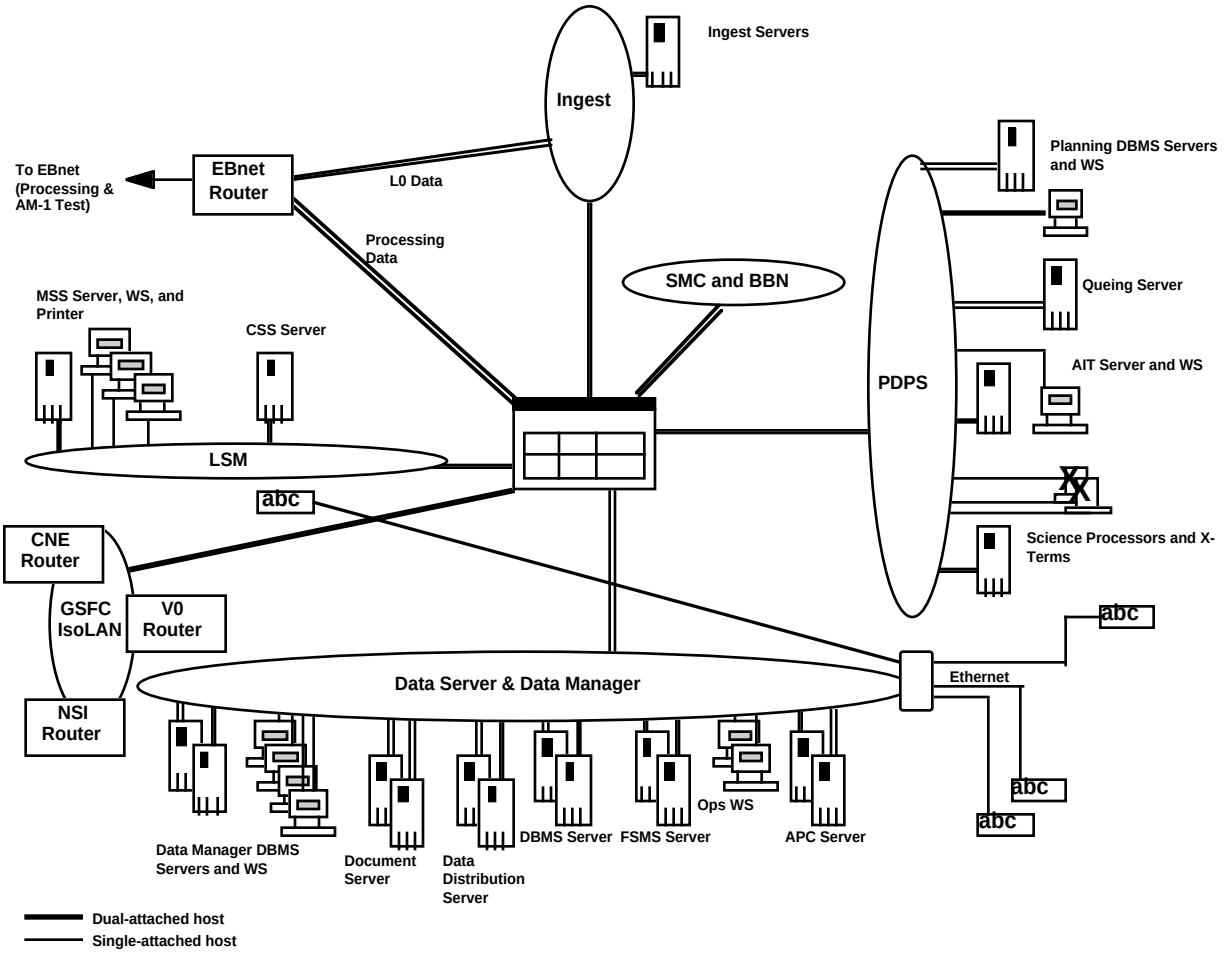

Figure 7.2.3.1-1 GSFC Release A DAAC LAN

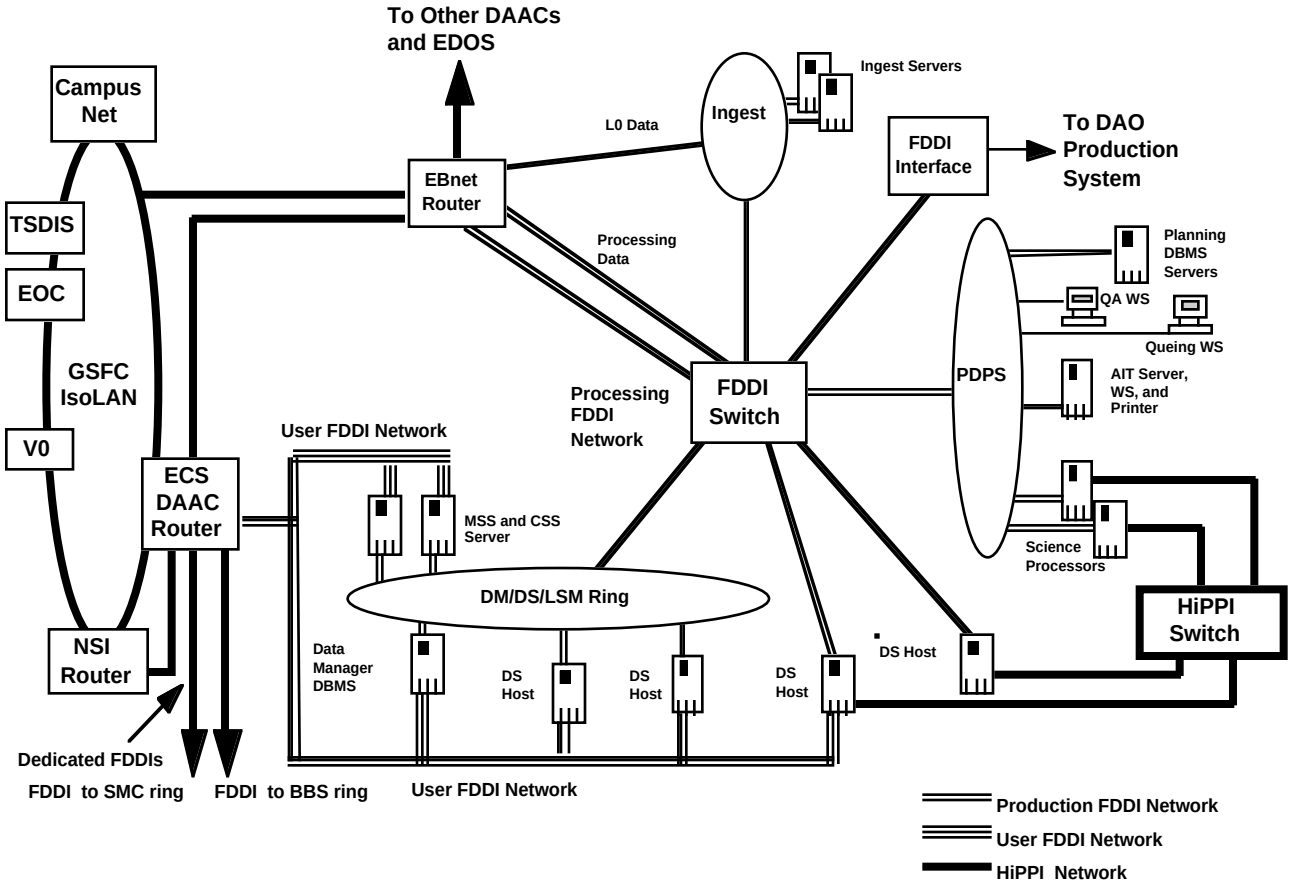

Figure 7.2.3.1-2 GSFC Release B DAAC LAN

7.2.10.2.2 LaRC ECS DAAC Network

Because LaRC DAAC's Release A and B network designs are quite similar to those of GSFC, the LaRC transitioning effort is very similar to that of GSFC. Figures 7.2.3.2-1 and 7.2.3.2-2 show the Release A and B designs, respectively. Significant differences between the LaRC and the GSFC transitioning are the following:

- The LaRC EBnet router is not connected directly to the campus IsoLAN nor is it connected to the ECS router.
- The EBnet router to FDDI Switch connection is not increased.

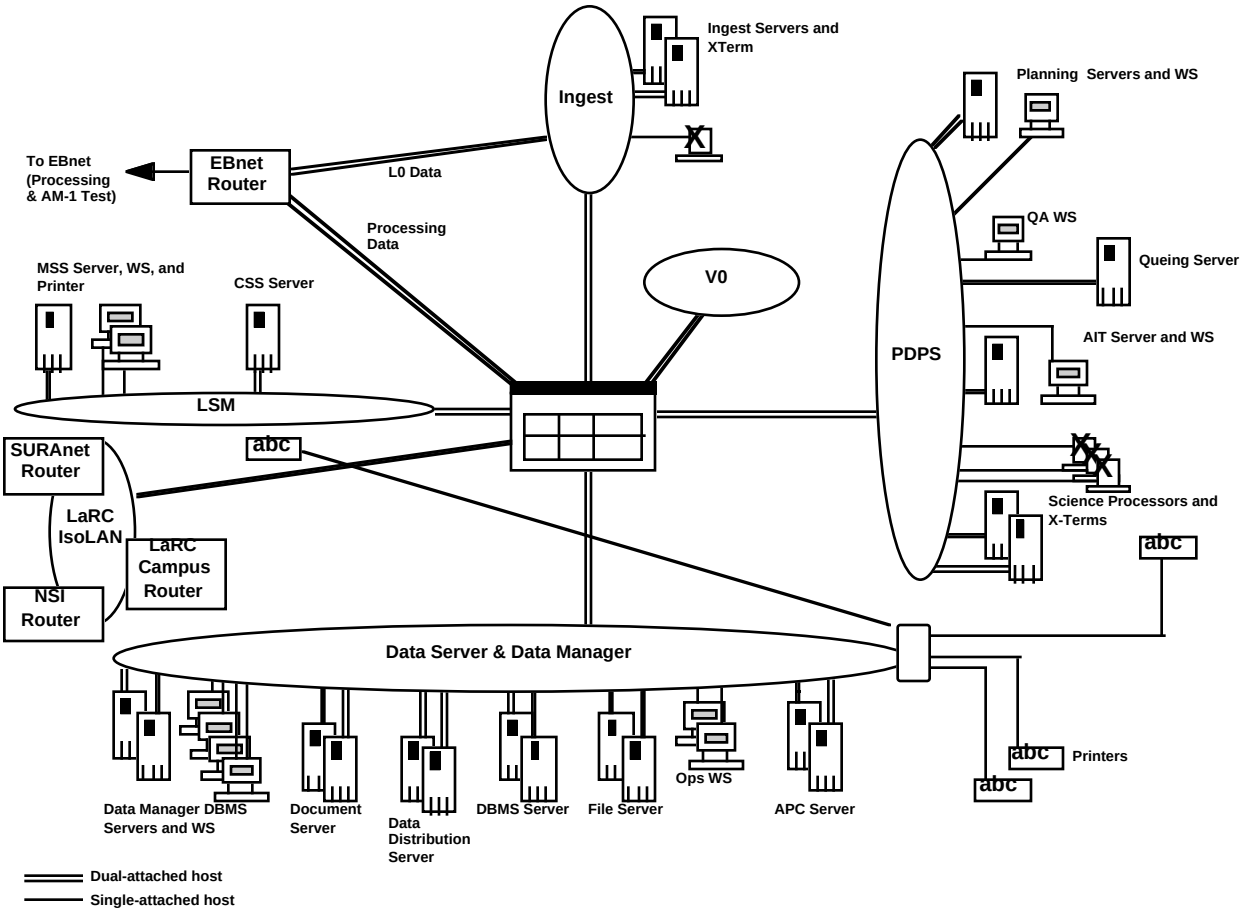

Figure 7.2.3.2-1 LaRC Release A DAAC LAN

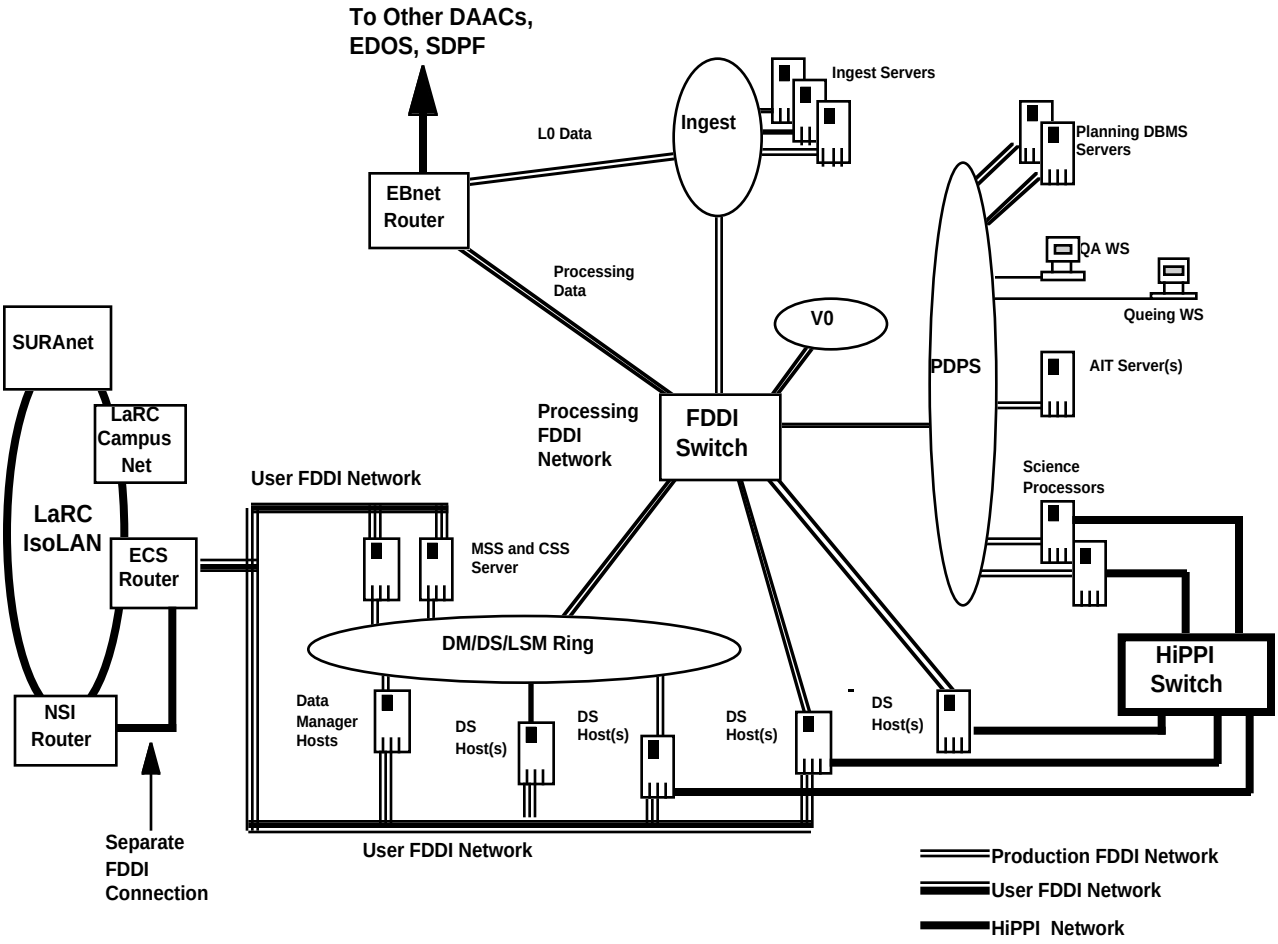

Figure 7.2.3.2-2 LaRC Release B DAAC LAN

7.2.10.2.3 EDC ECS DAAC Network

EDC's ECS DAAC network will undergo the most extensive change of the three DAACs because its Release A network configuration (as-built) is a simple, Ethernet-based topology with only a few subsystem hosts present. The subsystem hosts are all attached to an Ethernet hub that is part of the site's DAAC LAN (V0). Refer to Figure 7.2.3.3-1.

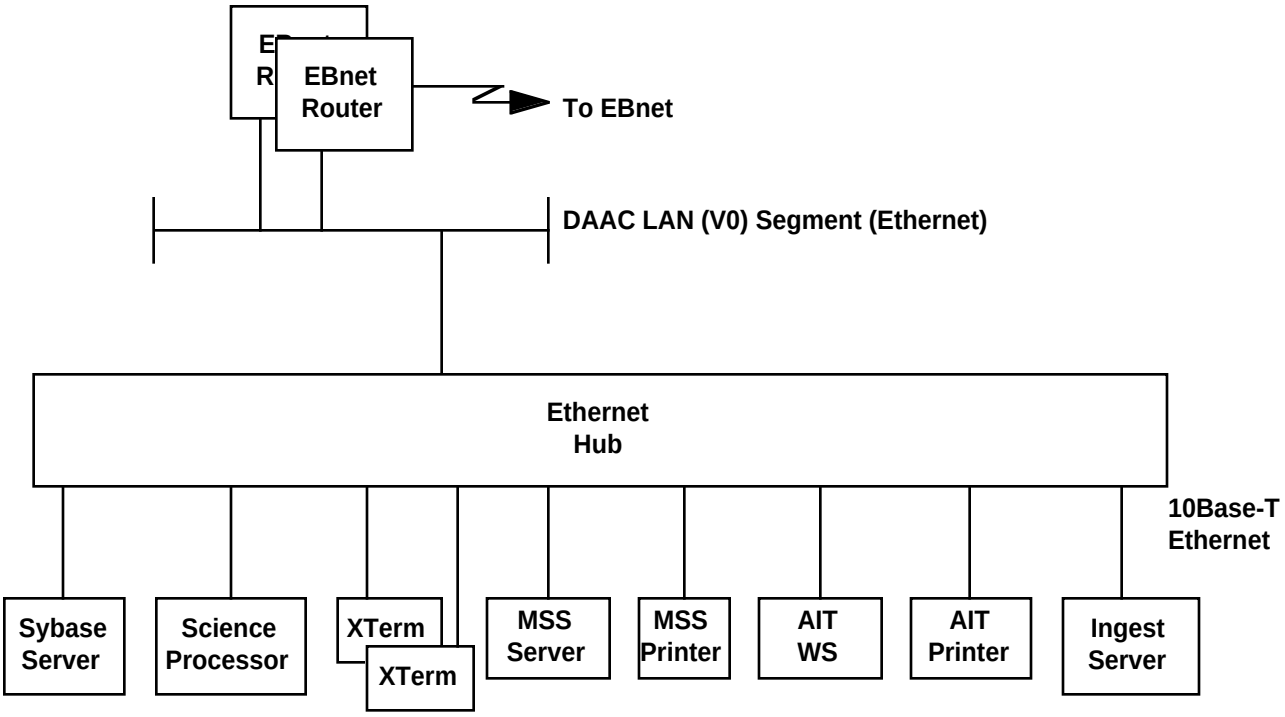

Figure 7.2.3.3-1 EDC Release A DAAC LAN

The new Release B network will be an entirely new physical network with its design consistent with the ECS generic network design. Figure 7.2.3.3-2 shows the Release B network for EDC. All Release A hosts retained for Release B will be moved to the new LAN.

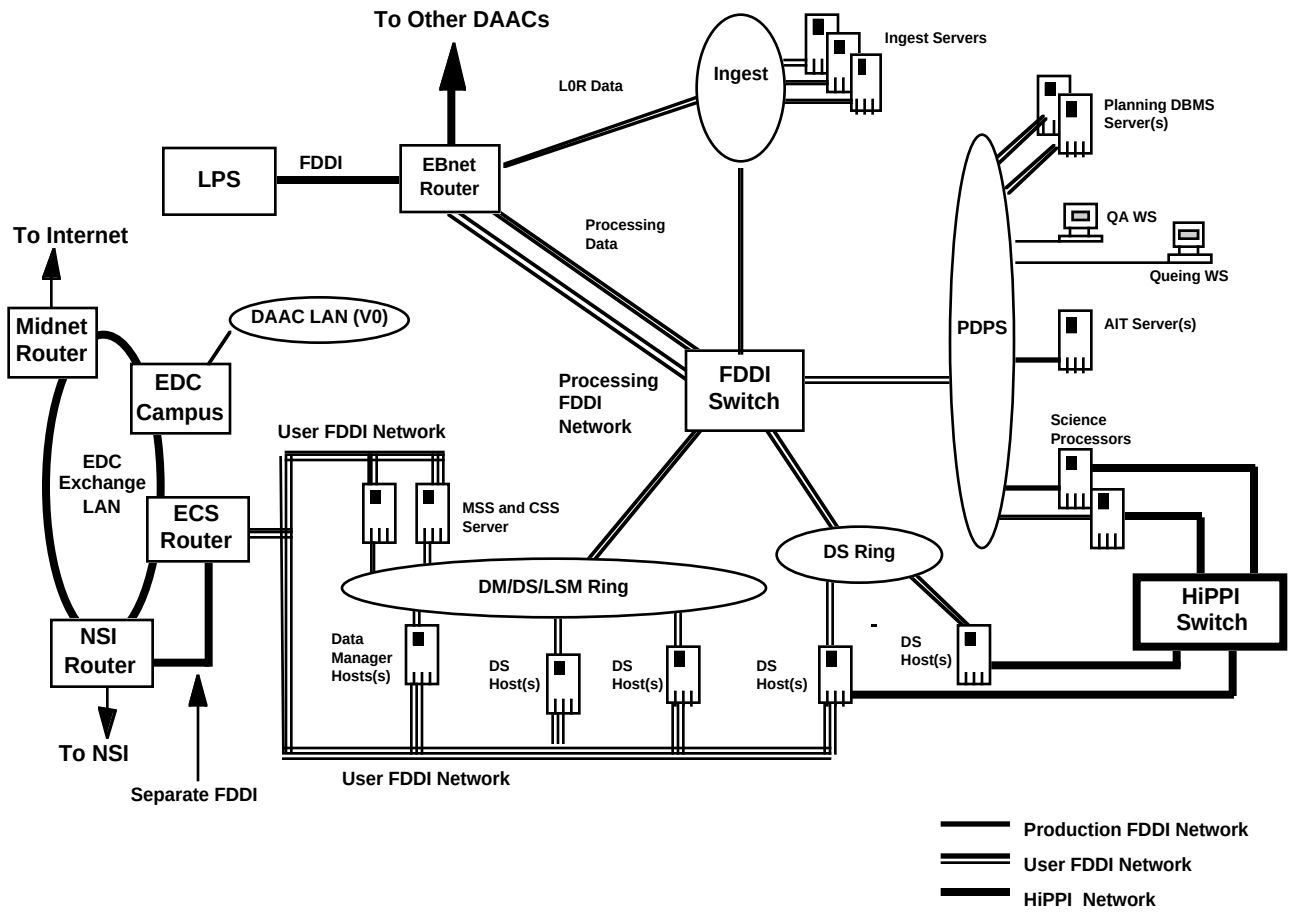

Figure 7.2.3.3-2 EDC Release B DAAC LAN

EDC does no Release A production. Therefore, the detailed transition plan will take into account any on-going Release A activity that must be on-going during the installation and test of Release B. In fact, there may be both networks present during the transition.

The new ECS Release B network and attached hosts will receive new IP network and host address assignments from EBnet. EBnet will modify the inter-DAAC bandwidth to accommodate Release B data volume.

The campus Exchange LAN (currently Ethernet-based) will be modified by EDC to add the ECS router for the connection to the Release B network. EDC will provide a FDDI interface to the Exchange LAN for this connection. In addition, the FDDI connection between the ECS router and the NSI router that will handle Release B user traffic will require the NSI router to be upgraded with a FDDI interface (this upgrade will be identified in an ICD for the ECS DAAC).

7.2.10.3 Assumptions and Dependencies

7.2.10.3.1 Assumptions

The following assumptions apply for the ISS network transitioning:

- Network routing protocol - RIP is retained for Release B intra-DAAC network routing. This assumption is the basis for the EBnet routing dependency below.
- Hardware availability - The general assumption that the Release B hardware is available at the sites includes the ISS networking hardware (e.g., router(s), hubs, concentrator cards, host interface cards, FDDI switch, and HiPPI switch)
- Down time - Down time (partial network and host) will be necessary for installation of host-specific network hardware interfaces and software device drivers (e.g., HiPPI NIC for the SGI platform, additional FDDI interface for MSS hosts), and for additions/changes to network hardware equipment, modifying configurations (e.g., router configuration), and changes to network router connections.
- Host network upgrades - Adequate host IO interface slots are available within hosts requiring additional and/or changed NIC interfaces from Release A to Release B networks. This would include, for example, one or more slots on a Science Processor for the HiPPI interface controller(s) and an IO interface slot on an MSS server for an additional FDDI interface card.
- Network traffic - Increases in DAAC LAN and EBnet WAN data flows resulting from transition to B testing will not impact on-going production. Estimate of Release B designed network capacity will be available. Anticipated over-capacity situations may be avoided by testing in non-production hours or subsystem-specific scheduling.
- BGP routing support - One or more DAAC exchange LANs may require the ECS router to support BGP routing protocol, thereby influencing the router selection.
- HiPPI technology deployed - The HiPPI network solution for Data Server-to-Science Processor subsystems' high data exchange rates is the selected solution.
- Air conditioning/power/facilities - Each DAAC will provide adequate air conditioning, power, and space to support the additional network equipment required for Release B.

7.2.10.3.2 Dependencies

The following dependencies apply for the ISS network transitioning:

- LAN addressing - Release A TCP/IP IP network and host addressing assignments made by ECS MRS Network Engineering will take into account the transition effort with the goal of minimizing changes for Release B at GSFC and LaRC. EDC's addressing will be changed. Further, network address assignments requested/received from EBnet will take into account the Release B design requirements.
- EBnet WAN:

- EBnet provides Release B operational routing resources, addressing assignment(s), and WAN connectivity and bandwidth that are consistent with the transitioning effort.
- EBnet DAAC site changes will be performed in a non-disruptive manner. EBnet's transition effort will be documented and coordinated with ECS transitional plans and implementation.
- EBnet may be required to simultaneously support Release A and Release B network connections at EDC.
- EBnet will continue to support RIP routing protocol at each ECS DAAC network interface for Release B. Inter-DAAC routing protocol is not relevant. ECS is to confirm this with EBnet and coordinate the plans and efforts with an ECS-requested EBnet transition plan.
- DAAC support - The transition effort is dependent upon coordinated support from DAAC LAN engineering, administration, and operations along with EBnet and M&O personnel for site network changes (e.g., facilities changes, adding the ECS router to the site's exchange LAN for Release B WAN connectivity).
- NSI router upgrade - EDC network modifications include an agreed upon upgrade to the NSI router. ECS will document the NSI support requirements in the appropriate ICD.

7.2.10.4 Data Changes

(Not applicable to the ISS Subsystem's transition.)

7.2.10.5 ISS Network Hardware Additions and Upgrades

7.2.10.5.1 GSFC DAAC

The following network hardware additions to the GSFC network are necessary:

- ECS router
- HiPPI switch with appropriate test equipment and configuration support equipment (VT100-type)
- FDDI and HiPPI interface cards
- FDDI concentrator(s) and/or FDDI ports in existing concentrators

The following network hardware upgrades to the GSFC network are necessary:

- EBnet router WAN bandwidth and number of FDDI connections
- Network Cables

7.2.10.5.2 LaRC DAAC

The following network hardware additions to the LaRC network are necessary:

- ECS router

- HiPPI switch with appropriate test equipment and configuration support equipment (VT100-type)
- FDDI and HiPPI interface cards
- FDDI concentrator(s) and/or FDDI ports in existing concentrators

The following network hardware upgrades to the LaRC network are necessary:

- EBnet router WAN bandwidth

7.2.10.5.3 EDC DAAC

The following network hardware additions to the EDC network are necessary:

- Full complement of Release B ISS network hardware (as-designed)

The following network hardware upgrades to the EDC network are necessary:

- EBnet router for additional WAN bandwidth and additional FDDI connections
- NSI router for FDDI connection to the ECS router

7.2.10.6 COTS Upgrades/Changes

(Not applicable to the ISS Subsystem's transition.)

7.2.10.7 Interoperability With Release A Components

1.≠ None of these changes will prevent Release A software (with minor changes, such as IP address changes) from executing with the Release B network in place.

2.≠ There will be no piecemeal changes to the network. All network changes at any one site will be completed at one time, to avoid testing multiple configurations.

7.2.10.8 Other Considerations

7.2.10.8.1 Bandwidth Availability

Transitioning Release A to Release B will require that the ECS network provide sufficient bandwidth to support concurrent Release A production and Release B testing. Therefore, the network transition effort will include measuring Release A data volume on portions of the network that will remain in the Release B network topology. This information along with Release B's data volume estimates and Release B's topology will be made available to aid the Release B subsystem transition's goal of avoiding disrupting Release A production.

7.2.10.8.2 Network Downtime

Network topology changes needed for Release B can be accomplished with minimal ECS Release A production downtime - but not without any downtime. Downtime estimates will be made part of the detailed transition implementation plan, along with mitigating options. Assuming that it is financially infeasible to duplicate each affected network and subsystem

hardware/software element, the following is a list of probable impacts on hosts and network elements that are shared in Release A and B:

- Hosts that receive a second network attachment (e.g., a second FDDI connection for an MSS server and a HiPPI network interface for a Data Server host) will have to be shut down in order to safely install the additional network interface card.
- FDDI concentrator(s) may be added to an existing FDDI ring, during which time the ring's traffic may be interrupted.
- Introducing the new ECS router into the topology will block the user communities access to the ECS DAAC during the installation and briefly following the installation completion as new network routes are propagated.

8. Candidate Solutions

8.1 Upgrade Release A Infrastructure To Match Release B

8.1.1 COTS

In order to minimize the perturbation to operations (Release A) during testing of Release B, it is highly desirable that the infrastructure (DCE version and topology, and operating system versions) be consistent between the two releases. Without a consistent infrastructure tests could not be run in parallel and would require substantial set-up and take-down time. Therefore the first phase of transition will be to install the Release B infrastructure into Release A, via the sustaining engineering process. This sustaining engineering release will be referred to as M&O.1. The following table defines the Release B versions of COTS that will need to be installed via M&O.1 to provide a common infrastructure for the two releases.

Table 8-1. COTS Versions For Common Infrastructure

Function	Product Name/ Vendor/ Features	Platforms Supported by Vendor	Vendor's Revision Level for REL A	Vendor's Revision Level for REL B
Operating Systems	Solaris	Sun	2.4	2.5
	HP-UX	HP	9.05 for w/s 10.0 for servers	10.0.1 (note DCE 1.1 needs 10.0)
	Digital UNIX (formerly OSF/1)	DEC	3.0/3.2 (FOS)	4.1
	SGI IRIX	SGI 32 bit SGI 64 bit	5.3 (32 bit) 6.1 (64 bit)	5.4 6.2
	AIX	IBM	3.2.5	4.1
Object Oriented DCE	OODCE (from Hewlett Packard) (Note: must upgrade with DCE 1.1) Client code must be recompiled.	HP	1.5.1	OSF/DCE 1.1
		Sun Solaris	1.5.1	OSF/DCE 1.1
		SGI		OSF/DCE 1.1
		IBM		OSF/DCE 1.1
		DEC		OSF/DCE 1.1
DCE Client and Server Software	from Transarc	Sun	1.0.3a	OSF/DCE 1.1
	from HP	HP	1.2.1	OSF/DCE 1.1
	from DEC	DEC	1.3	OSF/DCE 1.1
	from SGI	SGI	1.02	OSF/DCE 1.1
	from IBM	IBM	1.3 (& 1.2)	OSF/DCE 1.1

8.1.2 Mode Management Aspects

Where system components are shared at runtime by the 2 releases running in parallel under mode management, Release B software must conform to the Release A interface. These shared components include COTS (eg. DCE, Operating Systems, Sybase server) and multimode servers such as the archive server. Multimode servers are not part of Release A, and must be installed in the M&O.1 release.

8.1.3 DCE Cell Topology: Single Cell To Multi-Cell Transition

Any references to site in Release A code (in a CDS name for example, given that names can take the form /.:/ecs/$SITE/$MODE/...) will need to change for Release B since each site will become an independent cell. Neither names, nor the fact that site is an element of a name in Release A, should be hardcoded. A CDS wrapper will be delivered in Release A to hide the difference in naming between A and B (this wrapper is needed in any case to enforce mode management naming rules). The wrapper will inspect a configuration parameter to determine whether to use the Release A or B form of names. Use of this parameter will avoid the need to install new executables at the time that the change to multiple cells occurs.

Note that in DCE 1.0.3 cell names are exceedingly difficult to change once they are in place. The cell name is stored in many places within the cell, and perhaps outside the cell. For example the cell name is in every ACL entry in the cell, and in much of the CDS naming information. The cell name should not be changed after there has been any access to or from a foreign cell, since the old name will have propagated to the foreign cell. With DCE 1.1 aliasing provides a solution. Further study is needed to assess whether there are performance impacts.

9. Site Plans

This section provides one specific prototype transition strategy for a generic Release A DAAC, and for a generic Release B DAAC, making the assumption that the M&O.1 sustaining engineering release as described in section 8.1 is deployed.

Current plans are that Release B hardware is at sites by March 1997, while the software is pre∞ installed for early Release B walkthoughs in May- June 1997.

Later versions of this paper (see table 1-1) will include the sequence of changes for each site, together with schedules showing necessary downtime.

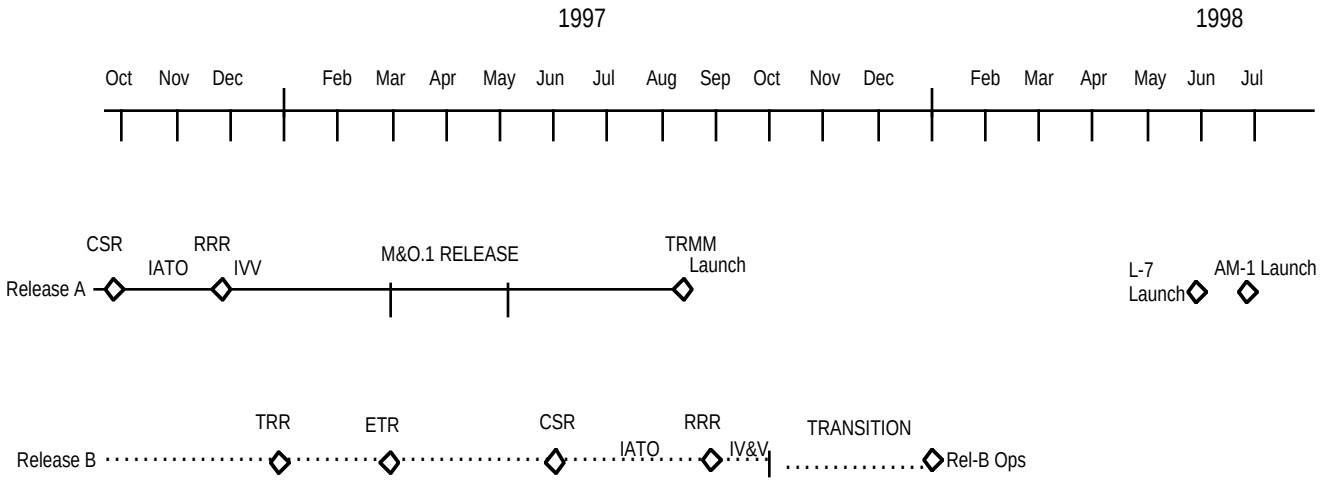

Figure 9-1. Timeframe for transitional activities

9.1 Release A Sites

The following series of steps transitions the Release A system to the Release B infrastructure (operating systems, DCE 1.1, and multicell topology) and installs the mode management capability, via the M&O.1 sustaining engineering release. This part of the transition will occur in spring 1997. The Release B system software arrives at sites for acceptance testing (IATO) and IV&V in summer and fall 1997, and transitional activities, such as migration of Release A data into Release B, occur in the last quarter of 1997.

9.1.1 Operating System Upgrades

Upgrades to the Release B operating system versions (defined in table 8-1) can occur processor by processor at each Release A site, to avoid shutting down the whole system. The 1.0.3 to 1.1

DCE prototyping effort, as a by-product, is intended to establish that Release A software runs on the Release B operating system versions. DCE 1.1 requires the Release B operating system versions.

9.1.2 DCE Upgrade

There are two options to upgrade to a new DCE version. The first option is to stop all machines in the cell (ie. to stop the entire Release A system), make the adjustments, and restart. The second option is to bring up a 1.1 cell in parallel with the Release A cell, and then configure each DCE client process to the new cell. This involves only a brief pause in operations for each DCE client. The 1.0.3 to 1.1 prototyping effort is proving the feasibility of operating OODCE 1.0.3 clients with 1.1 servers. The steps involved in this second option are as follows:

1. Bring up the DCE 1.1 name, security, and time servers on the CSS server processor, to establish the new cell. Since the master DCE servers are normally resident at the SMC, the SMC would be the default choice for location. To minimize performance impact and risk to the existing cell, the DCE 1.0.3 servers should be run on the MSS server so that the CSS server box is free to bring up the new cell. Copy the contents of the name and security databases from the 1.0.3 cell to the 1.1 cell.
2. Work through each site reconfiguring each DCE client from the old cell to the new cell, as follows:
   a. Unconfigure (ie. stop) the client.
   b. Change the file that specifies the location of the DCE name, security, and time servers.
   c. Restart the client.

   This is expected to take only a few minutes per box. Users can move to another workstation during this time.
3. At this stage all the 1.0.3 clients are operating in the 1.1 cell and the 1.0.3 cell can be shut down.
4. The 1.0.3 clients will be replaced with 1.1 clients in the M&O.1 release of custom code.

9.1.3 DCE Topology Change

This process is similar to the upgrade to DCE 1.1 since it involves the creation of a new cell - in this case at each site - and the reconfiguring of clients to the new cell. The series of steps is as follows:

1. At each Release A site (EDC, GSFC, LaRC), although not necessarily at the same time, bring up the DCE 1.1 name, security, and time servers to establish the new cell. As above, to minimize performance impact and risk to the existing cell, servers that usually run on the CSS server box should be moved to the MSS server.

2.≠ Copy segments for this site from the nameserver database in the old cell. Note that the structure of the CDS nameserver namespace changes with the topology, since each new cell replaces one branch of the old namespace.

3.≠ Populate the security database in the new cell, and remove corresponding entries in the old cell; this is not a simple copy operation because of security considerations with encrypted passwords, as well as format changes. All DCE users must be logged out while this step occurs.

4.≠ Work through the site reconfiguring each DCE client from the old cell to the new cell, as follows. As above, users can move to another workstation while the change is made.

   a. Unconfigure (ie. stop) the client.

   b.≠ Change the file that specifies the location of the DCE name, security, and time servers.

   c.≠ Change the configuration parameter that specifies to ECS processes whether to use single-cell or multi-cell names.

   d. Restart the client.

5. At this stage all DCE clients at the site are operating in the local cell.

6.≠ When GSFC DAAC, EDC, and LaRC have completed, then the SMC, assuming the master DCE servers for the old cell were resident there) can go through the same process to become its own independent cell, and the old cell can be shut down.

7. All Release A sites are now operating as independent cells.

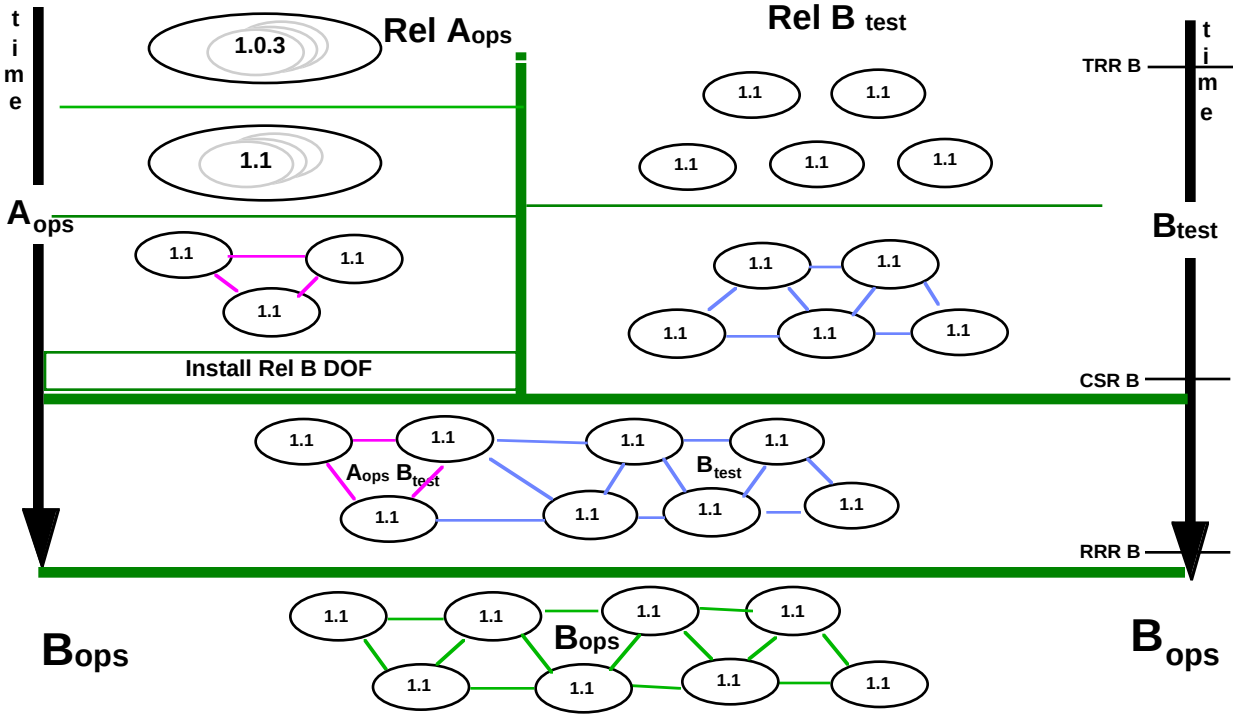

Figure 9-2 DCE Cell/Version Transition Approach

This figure shows the above sequence of steps, leading to completion of the DCE version upgrade and the cell topology change.

9.1.4 M&O.1 Sustaining Engineering Release

The M&O.1 release brings all DCE clients up to OODCE 1.1 and installs the mode management capability. Because mode management is not fully implemented in Release A (the infrastructure is in Release A but the full parallel operations and test capability is not) it will not be possible to test M&O.1 in parallel with A; however the size of the custom code changes in M&O.1 is relatively small, so operations outages due to test will not be extensive. Testing will focus on revalidation of operational paths in servers that have been upgraded to provide multimode capabilities.

9.1.5 Install Release B Hardware

The addition of the Release B hardware to the Release A system will occur prior to the IATO activities scheduled to start after CSR in June 1997. The network upgrade process for a release A site is as follows:

1.≠ Add ECS router to campus exchange (IsoLAN) network: Attach the router to the exchange LAN and configure it (including IP addresses and filter tables) for EBnet, SMC, BBN, and user FDDI network interfaces.

2.≠ Upgrade EBnet: Modify EBnet WAN bandwidth and/or WAN connectivity, configure the EBnet router to support additional WAN bandwidth, and configure the EBnet router to support additional FDDI interfaces.

3.≠ Prepare HiPPI switch and host cabling: Install and configure the HiPPI switch, position HiPPI cabling between the HiPPI switch and host locations, test the HiPPI switch in loop-back mode, and attach cabling to it.

4.≠ Upgrade FDDI switch by adding required interfaces for ECS networks. Configure the switch for additional ports and modify the filter table.

5.≠ Construct user FDDI network: Install FDDI concentrator unit(s), lay FDDI cabling between concentrators and host(s), attach concentrators to ECS router.

6.≠ Add second network interface card to hosts: add HiPPI card and device driver software to SGI hosts in Data Server and Science Processors, and add FDDI card to LSM host(s) and Data Server host(s). Add new IP addresses to host configuration.

7.≠ Complete the HiPPI network installation: for each host, configure the HiPPI interface (includes new IP address), test the host HiPPI interface with a loop-back cable, connect HiPPI cabling from HiPPI switch to host.

8.≠ Complete user network construction by attaching hosts to the user FDDI network and configure the interfaces with the new IP addresses. Existing connections' IP addresses do not change.

9.≠ Move LSM-attached hosts and concentrators to DM/DS/LSM FDDI network. IP addresses do not change.

10.≠Complete Data Server-to-FDDI host connections by attaching Data Server hosts to the FDDI switch and configure the interfaces with the new IP addresses. IP addresses of existing connections do not change.

11.≠Complete the SMC connection by moving the SMC host FDDI concentrator connection to the ECS router interface. SMC host IP addresses do not change.

12.≠Complete the BBN connection by moving the BB FDDI concentrator to the ECS router interface. Modify the host configuration with the changed IP address.

9.1.6 IATO And IV&V For Release B

IATO activities will begin at Release-B only sites and move on to Release A sites when experience is gained with both the software under test and the test procedures. At the Release A sites, acceptance testing and IV&V can be performed in parallel with operations using mode management, to the extent that they are compatible; stress testing and failover testing will not be compatible. After checkout at each site, IATO activities include a full-up test with all sites participating.

Note that by default the SMC will become aware of the new Release B-only sites as they come online, due to the IP discovery feature of HP Openview. Optionally, Openview could be configured to ignore non-operational Release B sites to avoid confusion with Release A operations. The SMC could take advantage of mode management to operate a test mode that coordinated Release B systems, while the operational mode coordinated Release A systems.

9.1.7 Transitional Activities

Following IATO and IV&V the Release B system is verified and validated. The remaining activities prior to operations include:

1.≠ Folding of data from Release A databases into Release B databases, as discussed for each subsystem in section 7.2:

   MSS management database (Sybase V10 to V11)

   PDPS database (schema changes and Sybase upgrade)

   Science Data Server (Sybase to Illustra migration)

   Document Data Server (movement of keyword index and documents)

   DDICT (migration of data from GTWAY)

   Ingest database (Sybase V10 to V11)

2.≠ Execution of Releases A and B with the same data inputs and persistent state data, to verify that identical results are produced.

9.2 Release B-Only Sites

Transition to B at the sites that did not have Release A installed is the same as installation, with the following exceptions:

1.≠ Release B will be tested at the least busy B-only sites prior to the Release A sites in order to gain experience and shakedown the system thoroughly before attempting the more complex transitions at the Release A sites.

2.≠ There are existing systems at some Release B-only sites that will be impacted by the introduction of ECS. ASF and ORNL, at least, will be impacted by the replacement of parts of their existing system by ECS-B.

Abbreviations and Acronyms

ADSRV	advertising service CSCI
AHWGP	Ad Hoc Working Group for Production
AI&T	Algorithm Integration and Test
AMASS	archival management and storage system
APC	Access/Process Coordinators
API	application program interface
AS	Administration Stations
ASCII	American Standard Code for Information Interchange
ATL	automatic tape library
ATM	asynchronous transmission mode
CD-ROM	compact disk - read only memory
CDR	Critical Design Review
CDS	Cell Directory Service
CGI	Common Gateway Interface
CI	Configuration Item
CLS	Client Subsystems
CORBA	Common Object Request Broker Architecture
COTS	Commercial off-the Shelf
CPU	you should know this one
CSC	Computer Software Component
CSCI	Computer Software Configuration Item
CSMS	Communications and System Management Segment
CSS	Communications And Infrastructure Subsystem
DAAC	Distributed Active Archive Center
DAS	Dual-Attached Station
DBA	database administrator
DBMS	database management system
DCE	distributed computing environment
DDICT	data dictionary CSCI
DDIST	data distribution CSCI

DDSRV	document data server CSCI
DEV	developed code
DMS	Data Management Subsystem
DPR	Data Processing Request
DSS	Data Server Subsystem
ECS	EOSDIS Core System
EDF	ECS development facility
EOSDIS	Earth Observation System Data Information System
ESDT	Earth science data typesECS EOSDIS Core System
ESO	ECS Server Object
FOS	Flight Operations Segment
FSMS	File Storage Management System
FTP	file transfer protocol
GDS	Global Directory Service
GSFC	Goddard Space Flight Center
GSO	Global Server Object
HP	Hewlett Packard
HTML	hypertext markup language
HTTP	hypertext transfer protocol
I&T	Integration and Testing
INGST	Ingest CI
INS	Ingest Subsystem
IOS	Interoperability Subsystem
ISS	Internetworking Subsystem
MDA	Management Data Access
MIB	Management Information Base
MPF	Managed Process Framework
MMS	Mode Management Service
MSS	Management Services Subsystem
NASA	National Aeronautics and Space Administration
NFS	Network File System
OODCE	Object Oriented DCE
OPS	operations

OTS	off-the-shelf
PCF	Process Control File
PDPS	Planning and Data Processing Subsystem
PDR	Preliminary Design Review
PGE	Product Generation Executable
PGS	Product Generation System
PLS	Planning Subsystem
PRONG	Processing CI
PRS	Processing Subsystem
RAID	redundant array of inexpensive disks
RMA	reliability, maintainability, availabilityRDA Remote Data Access
RPC	Remote Procedure Call
SCF	Science Computing Facility
SDP	Science Data Processing
SDPS	Science Data Processing Segment
SDSRV	Science Data Server CSCI
SMC	System Monitoring and Coordination
SMP	symmetric multi-processor
SNMP	Simple Network Management Protocol
SSI&T	Science Software I&T
SSJ	Server Session Job
STMGT	Storage Management Software CSCI
TRMM	Tropical Rainfall Measuring Mission
TSDIS	TRMM Science Data and Information System
URDB	User Registration DB
URL	Uniform Resource Locators
V0	Version 0
V1	Version 1

This page intentionally left blank.